

# AltiVec™ Technology Programming Environments Manual




# Freescale Semiconductor, Inc.

AltiVec is a trademark of Motorola, Inc.  
DigitalDNA is a trademark of Motorola, Inc.

The PowerPC name and the PowerPC logotype are trademarks of International Business Machines Corporation used by Motorola under license from International Business Machines Corporation.

This document contains information on a new product under development. Motorola reserves the right to change or discontinue this product without notice. Information in this document is provided solely to enable system and software implementers to use PowerPC microprocessors. There are no express or implied copyright licenses granted hereunder to design or fabricate PowerPC integrated circuits or integrated circuits based on the information in this document.

Motorola reserves the right to make changes without further notice to any products herein. Motorola makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Motorola assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters can and do vary in different applications. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Motorola does not convey any license under its patent rights nor the rights of others. Motorola products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Motorola product could create a situation where personal injury or death may occur. Should Buyer purchase or use Motorola products for any such unintended or unauthorized application, Buyer shall indemnify and hold Motorola and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Motorola was negligent regarding the design or manufacture of the part. Motorola and  are registered trademarks of Motorola, Inc. Motorola, Inc. is an Equal Opportunity/Affirmative Action Employer.

#### Motorola Literature Distribution Centers:

**USA/EUROPE:** Motorola Literature Distribution; P.O. Box 5405; Denver, Colorado 80217; Tel.: 1-800-441-2447 or 1-303-675-2140/

**JAPAN:** Nippon Motorola Ltd SPD, Strategic Planning Office 4-32-1, Nishi-Gotanda Shinagawa-ku, Tokyo 141, Japan Tel.: 81-3-5487-8488

**ASIA/PACIFIC:** Motorola Semiconductors H.K. Ltd.; 8B Tai Ping Industrial Park, 51 Ting Kok Road, Tai Po, N.T., Hong Kong; Tel.: 852-26629298

**World Wide Web Address:** <http://sps.motorola.com/mfax>

**INTERNET:** <http://motorola.com/sps>

**Technical Information:** Motorola Inc. SPS Customer Support Center 1-800-521-6274; electronic mail address: [crc@wmkmail.sps.mot.com](mailto:crc@wmkmail.sps.mot.com).

**Document Comments:** FAX (512) 933-2625, Attn: RISC Applications Engineering.

**World Wide Web Addresses:** <http://www.mot.com/PowerPC>  
<http://www.mot.com/netcomm>

© Motorola Inc. 2001. All rights reserved.

**For More Information On This Product,  
Go to: [www.freescale.com](http://www.freescale.com)**

# Freescale Semiconductor, Inc.

Overview	1
AltiVec Register Set	2
Operand Conventions	3
Addressing Modes and Instruction Set Summary	4
Cache, Exceptions, and Memory Management	5
AltiVec Instructions	6
Appendix A: Instruction Set Mnemonics - Decimal	A
Appendix B: Instruction Set Mnemonics - Binary	B
Appendix C: Opcodes - Decimal	C
Appendix D: Opcodes - Binary	D
Appendix E: Forms	E
Appendix F: Legends	F
Appendix G: Revision History	G
Glossary of Terms and Abbreviations	GLO
Index	IND

# Freescale Semiconductor, Inc.

1

Overview

2

AltiVec Register Set

3

Operand Conventions

4

Addressing Modes and Instruction Set Summary

5

Cache, Exceptions, and Memory Management

6

AltiVec Instructions

A

Appendix A: Instruction Set Mnemonics - Decimal

B

Appendix B: Instruction Set Mnemonics - Binary

C

Appendix C: Opcodes - Decimal

D

Appendix D: Opcodes - Binary

E

Appendix E: Forms

F

Appendix F: Legends

G

Appendix G: Revision History

GLO

Glossary of Terms and Abbreviations

IND

Index

## Contents

Paragraph Number	Title	Page Number
	Audience .....	xx
	Organization.....	xxi
	Suggested Reading.....	xxi
	General Information.....	xxii
	Related Documentation .....	xxii
	Conventions .....	xxiii
	Acronyms and Abbreviations.....	xxiv
	Terminology Conventions.....	xxvii

### Chapter 1 Overview

1.1	Overview .....	1-1
1.2	AltiVec Technology Overview .....	1-3
1.2.1	Levels of AltiVec ISA .....	1-5
1.2.2	Features Not Defined by AltiVec ISA.....	1-6
1.3	AltiVec Architectural Model.....	1-6
1.3.1	AltiVec Registers and Programming Model .....	1-6
1.3.2	Operand Conventions.....	1-7
1.3.2.1	Byte Ordering .....	1-7
1.3.2.2	Floating-Point Conventions .....	1-8
1.3.3	AltiVec Addressing Modes .....	1-9
1.3.4	AltiVec Instruction Set.....	1-11
1.3.5	AltiVec Cache Model .....	1-12
1.3.6	AltiVec Exception Model.....	1-12
1.3.7	Memory Management Model .....	1-12

### Chapter 2 AltiVec Register Set

2.1	Overview on the AltiVec and PowerPC Registers .....	2-1
2.2	AltiVec Register Set Overview .....	2-3
2.3	Registers defined by AltiVec ISA .....	2-4
2.3.1	AltiVec Vector Register File (VRF).....	2-4
2.3.2	Vector Status and Control Register (VSCR).....	2-4
2.3.3	Vector Save/Restore Register (VRSAVE).....	2-6

# Contents

Paragraph Number	Title	Page Number
2.4	Additions to PowerPC UISA Registers .....	2-7
2.4.1	PowerPC Condition Register .....	2-8
2.5	Additions to PowerPC OEA Registers.....	2-9
2.5.1	AltiVec Field added in the PowerPC Machine State Register (MSR) .....	2-9
2.5.2	Machine Status Save/Restore Registers (SRRs) .....	2-10
2.5.2.1	Machine Status Save/Restore Register 0 (SRR0) .....	2-10
2.5.2.2	Machine Status Save/Restore Register 1 (SRR1) .....	2-11

## Chapter 3 Operand Conventions

3.1	Data Organization in Memory .....	3-1
3.1.1	Aligned and Misaligned Accesses .....	3-1
3.1.2	AltiVec Byte Ordering .....	3-2
3.1.2.1	Big-Endian Byte Ordering .....	3-3
3.1.2.2	Little-Endian Byte Ordering .....	3-3
3.1.3	Quad Word Byte Ordering Example .....	3-3
3.1.4	Aligned Scalars in Little-Endian Mode .....	3-4
3.1.5	Vector Register and Memory Access Alignment .....	3-6
3.1.6	Quad-Word Data Alignment .....	3-7
3.1.6.1	Accessing a Misaligned Quad Word in Big-Endian Mode .....	3-8
3.1.6.2	Accessing a Misaligned Quad Word in Little-Endian Mode .....	3-10
3.1.6.3	Scalar Loads and Stores .....	3-11
3.1.6.4	Misaligned Scalar Loads and Stores .....	3-11
3.1.7	Mixed-Endian Systems .....	3-12
3.2	AltiVec Floating-Point Instructions—UISA .....	3-12
3.2.1	Floating-Point Modes .....	3-13
3.2.1.1	Java Mode .....	3-13
3.2.1.2	Non-Java Mode .....	3-14
3.2.2	Floating-Point Infinities .....	3-14
3.2.3	Floating-Point Rounding.....	3-14
3.2.4	Floating-Point Exceptions.....	3-14
3.2.4.1	NaN Operand Exception .....	3-15
3.2.4.2	Invalid Operation Exception .....	3-16
3.2.4.3	Zero Divide Exception .....	3-16
3.2.4.4	Log of Zero Exception .....	3-16
3.2.4.5	Overflow Exception .....	3-17
3.2.4.6	Underflow Exception .....	3-17
3.2.5	Floating-Point NaNs .....	3-17
3.2.5.1	NaN Precedence.....	3-18
3.2.5.2	SNaN Arithmetic .....	3-18
3.2.5.3	QNaN Arithmetic .....	3-18

# Contents

Paragraph Number	Title	Page Number
3.2.5.4	NaN Conversion to Integer .....	3-18
3.2.5.5	NaN Production .....	3-18

## Chapter 4 Addressing Modes and Instruction Set Summary

4.1	Conventions .....	4-2
4.1.1	Execution Model .....	4-2
4.1.2	Computation Modes .....	4-2
4.1.3	Classes of Instructions .....	4-2
4.1.4	Memory Addressing .....	4-3
4.1.4.1	Memory Operands .....	4-3
4.1.4.2	Effective Address Calculation .....	4-3
4.2	AltiVec UISA Instructions .....	4-4
4.2.1	Vector Integer Instructions .....	4-4
4.2.1.1	Saturation Detection .....	4-4
4.2.1.2	Vector Integer Arithmetic Instructions .....	4-5
4.2.1.3	Vector Integer Compare Instructions .....	4-13
4.2.1.4	Vector Integer Logical Instructions .....	4-15
4.2.1.5	Vector Integer Rotate and Shift Instructions .....	4-16
4.2.2	Vector Floating-Point Instructions .....	4-17
4.2.2.1	Floating-Point Division and Square-Root .....	4-18
4.2.2.1.1	Floating-Point Division .....	4-18
4.2.2.1.2	Floating-Point Square-Root .....	4-19
4.2.2.2	Floating-Point Arithmetic Instructions .....	4-19
4.2.2.3	Floating-Point Multiply-Add Instructions .....	4-20
4.2.2.4	Floating-Point Rounding and Conversion Instructions .....	4-21
4.2.2.5	Floating-Point Compare Instructions .....	4-22
4.2.2.6	Floating-Point Estimate Instructions .....	4-24
4.2.3	Load and Store Instructions .....	4-25
4.2.3.1	Alignment .....	4-26
4.2.3.2	Load and Store Address Generation .....	4-26
4.2.3.3	Vector Load Instructions .....	4-27
4.2.3.4	Vector Store Instructions .....	4-30
4.2.4	Control Flow .....	4-31
4.2.5	Vector Permutation and Formatting Instructions .....	4-31
4.2.5.1	Vector Pack Instructions .....	4-31
4.2.5.2	Vector Unpack Instructions .....	4-33
4.2.5.3	Vector Merge Instructions .....	4-34
4.2.5.4	Vector Splat Instructions .....	4-35
4.2.5.5	Vector Permute Instruction .....	4-36
4.2.5.6	Vector Select Instruction .....	4-36

## Contents

4.2.5.7	Vector Shift Instructions .....	4-37
4.2.5.7.1	Immediate Interelement Shifts/Rotates.....	4-37
4.2.5.7.2	Computed Interelement Shifts/Rotates .....	4-38
4.2.5.7.3	Variable Interelement Shifts .....	4-39
4.2.6	Processor Control Instructions—UISA .....	4-39
4.2.6.1	AltiVec Status and Control Register Instructions .....	4-40
4.2.7	Recommended Simplified Mnemonics .....	4-40
4.3	AltiVec VEA Instructions .....	4-40
4.3.1	Memory Control Instructions—VEA .....	4-41
4.3.2	User-Level Cache Instructions—VEA.....	4-41

### Chapter 5

#### Cache, Exceptions, and Memory Management

5.1	PowerPC Shared Memory.....	5-1
5.2	AltiVec Memory Bandwidth Management .....	5-1
5.2.1	Software-Directed Prefetch.....	5-2
5.2.1.1	Data Stream Touch ( <b>dst</b> ).....	5-2
5.2.1.2	Transient Streams .....	5-4
5.2.1.3	Storing to Streams ( <b>dstst</b> ).....	5-4
5.2.1.4	Stopping Streams .....	5-5
5.2.1.5	Exception Behavior of Prefetch Streams .....	5-6
5.2.1.6	Synchronization Behavior of Streams .....	5-7
5.2.1.7	Address Translation for Streams.....	5-7
5.2.1.8	Stream Usage Notes.....	5-7
5.2.1.9	Stream Implementation Assumptions .....	5-9
5.2.2	Prioritizing Cache Block Replacement.....	5-9
5.2.3	Partially Executed AltiVec Instructions .....	5-10
5.3	DSI Exception—Data Address Breakpoint.....	5-10
5.4	AltiVec Unavailable Exception (0x00F20).....	5-10

### Chapter 6

#### AltiVec Instructions

6.1	Instruction Formats .....	6-1
6.1.1	Instruction Fields .....	6-1
6.1.2	Notation and Conventions.....	6-2
6.2	AltiVec Instruction Set.....	6-8



# Contents

Paragraph Number	Title	Page Number
<b>Appendix A</b> <b>AltiVec Instruction Set Listings</b>		
A.1	Instructions Sorted by Mnemonic in Decimal Format.....	A-1
<b>Appendix B</b> <b>Instructions Sorted by Mnemonic in Binary Format</b>		
B.1	Instructions Sorted by Mnemonic in Binary Format .....	B-1
<b>Appendix C</b> <b>Instructions Sorted by Opcode</b>		
C.1	Instructions Sorted by Opcode in Decimal Format.....	C-1
<b>Appendix D</b> <b>Instructions Sorted by Opcode</b>		
D.1	Instructions Sorted by Opcode in Binary Format .....	D-1
<b>Appendix E</b> <b>Instructions Sorted by Form</b>		
E.1	Instructions Sorted by Form.....	E-1
<b>Appendix F</b> <b>Instruction Set Legend</b>		
F.1	Instruction Set Legend .....	F-1
<b>Appendix G</b> <b>User's Manual Revision History</b>		
G.1	Revision History .....	G-1

## Glossary

## Index

# Contents

**Paragraph  
Number**

**Title**

**Page  
Number**

**Freescale Semiconductor, Inc.**

## Figures

Figure Number	Title	Page Number
1-1	Overview of PowerPC architecture with AltiVec Technology .....	1-4
1-2	AltiVec Top-Level Diagram .....	1-7
1-3	Big-Endian Byte Ordering for a Vector Register .....	1-8
1-4	Bit Ordering .....	1-8
1-5	Intraelement Example, vaddsbs .....	1-9
1-6	Interelement Example, vperm .....	1-9
2-1	Programming Model—All Registers .....	2-2
2-2	AltiVec Register Set .....	2-3
2-3	Vector Registers (VRs).....	2-4
2-4	Vector Status and Control Register (VSCR) .....	2-5
2-5	32-bit VSCR Moved to a 128-bit Vector Register .....	2-5
2-6	Vector Save/Restore Register (VRSAVE).....	2-7
2-7	Condition Register (CR) .....	2-8
2-8	Machine State Register (MSR) .....	2-9
2-9	Machine Status Save/Restore Register 0 (SRR0) .....	2-11
2-10	Machine Status Save/Restore Register 0 (SRR1) .....	2-11
3-1	Big-Endian Mapping of a Quad Word .....	3-3
3-2	Little-Endian Mapping of a Quad Word .....	3-4
3-3	Little-Endian Mapping of Quad Word—Alternate View .....	3-4
3-4	Quad Word Load with PowerPC Munged Little-Endian Applied.....	3-5
3-5	AltiVec Little Endian Double-Word Swap.....	3-6
3-6	Misaligned Vector in Big-Endian Mode.....	3-7
3-8	Big-Endian Quad Word Alignment .....	3-8
3-7	Misaligned Vector in Little-Endian Addressing Mode.....	3-8
3-9	Little-Endian Alignment .....	3-11
4-1	Register Indirect with Index Addressing for Loads/Stores .....	4-27
5-1	Format of rB in dst Instruction.....	5-2
5-2	Data Stream Touch .....	5-3
5-3	SRR1 Bit Settings after an AltiVec Unavailable Exception .....	5-11
6-1	Format of rB in dst instruction (32-bit).....	6-13
6-2	Effects of Example Load/Store Instructions .....	6-15
6-3	Load Vector for Shift Left .....	6-18
6-4	Instruction vperm Used in Aligning Data .....	6-19
6-5	vaddcuw—Determine Carries of Four Unsigned Integer Adds (32-Bit) .....	6-30
6-6	vaddfp—Add Four Floating-Point Elements (32-Bit) .....	6-31

# Figures

Figure Number	Title	Page Number
6-7	vaddsbs—Add Saturating Sixteen Signed Integer Elements (8-Bit) .....	6-32
6-8	vaddshs—Add Saturating Eight Signed Integer Elements (16-Bit) .....	6-33
6-9	vaddsws—Add Saturating Four Signed Integer Elements (32-Bit) .....	6-34
6-10	vaddubm—Add Sixteen Integer Elements (8-Bit) .....	6-35
6-11	vaddubs—Add Saturating Sixteen Unsigned Integer Elements (8-Bit) .....	6-36
6-12	vadduhm—Add Eight Integer Elements (16-Bit) .....	6-37
6-13	vadduhs—Add Saturating Eight Unsigned Integer Elements (16-Bit) .....	6-38
6-14	vadduwm—Add Four Integer Elements (32-Bit) .....	6-39
6-15	vadduws—Add Saturating Four Unsigned Integer Elements (32-Bit) .....	6-40
6-16	vand—Logical Bitwise AND .....	6-41
6-17	vand—Logical Bitwise AND with Complement .....	6-42
6-18	vavgsh—Average Sixteen Signed Integer Elements (8-Bit) .....	6-43
6-19	vavgsh—Average Eight Signed Integer Elements (16-bits) .....	6-44
6-20	vavgsw—Average Four Signed Integer Elements (32-Bit) .....	6-45
6-21	vavgub—Average Sixteen Unsigned Integer Elements (8-bits) .....	6-46
6-22	vavgsh—Average Eight Signed Integer Elements (16-Bit) .....	6-47
6-23	vavguw—Average Four Unsigned Integer Elements (32-Bit) .....	6-48
6-24	vcfsx—Convert Four Signed Integer Elements to Four Floating-Point Elements (32-Bit) .....	6-49
6-25	vcfux—Convert Four Unsigned Integer Elements to Four Floating-Point Elements (32-Bit) .....	6-50
6-26	vcmpbfp—Compare Bounds of Four Floating-Point Elements (32-Bit) .....	6-52
6-27	vcmpeqfp—Compare Equal of Four Floating-Point Elements (32-Bit) .....	6-53
6-28	vcmpequb—Compare Equal of Sixteen Integer Elements (8-bits) .....	6-54
6-29	vcmpequh—Compare Equal of Eight Integer Elements (16-Bit) .....	6-55
6-30	vcmpequw—Compare Equal of Four Integer Elements (32-Bit) .....	6-56
6-31	vcmpgef—Compare Greater-Than-or-Equal of Four Floating-Point Elements (32-Bit) .....	6-57
6-32	vcmpgtfp—Compare Greater-Than of Four Floating-Point Elements (32-Bit) .....	6-58
6-33	vcmpgtsh—Compare Greater-Than of Sixteen Signed Integer Elements (8-Bit) .....	6-59
6-34	vcmpgtsh—Compare Greater-Than of Eight Signed Integer Elements (16-Bit) .....	6-60
6-35	vcmpgtsw—Compare Greater-Than of Four Signed Integer Elements (32-Bit) .....	6-61
6-36	vcmpgtub—Compare Greater-Than of Sixteen Unsigned Integer Elements (8-Bit) .....	6-62
6-37	vcmpgtuh—Compare Greater-Than of Eight Unsigned Integer Elements (16-Bit) .....	6-63
6-38	vcmpgtuw—Compare Greater-Than of Four Unsigned Integer Elements (32-Bit) .....	6-64
6-39	vctxs—Convert Four Floating-Point Elements to Four Signed Integer Elements (32-Bit) .....	6-65
6-40	vctxu—Convert Four Floating-Point Elements to Four Unsigned Integer Elements (32-Bit) .....	6-66
6-41	vxptefp—2 Raised to the Exponent Estimate Floating-Point for Four Floating-Point Elements (32-Bit) .....	6-68

# Figures

Figure Number	Title	Page Number
6-42	vexptefp—Log2 Estimate Floating-Point for Four Floating-Point Elements (32-Bit) .....	6-70
6-43	vmaddfp—Multiply-Add Four Floating-Point Elements (32-Bit) .....	6-71
6-44	vmaxfp—Maximum of Four Floating-Point Elements (32-Bit) .....	6-72
6-45	vmaxsb—Maximum of Sixteen Signed Integer Elements (8-Bit) .....	6-73
6-46	vmaxsh—Maximum of Eight Signed Integer Elements (16-Bit) .....	6-74
6-47	vmaxsw—Maximum of Four Signed Integer Elements (32-Bit) .....	6-75
6-48	vmaxub—Maximum of Sixteen Unsigned Integer Elements (8-Bit) .....	6-76
6-49	vmaxuh—Maximum of Eight Unsigned Integer Elements (16-Bit) .....	6-77
6-50	vmaxuw—Maximum of Four Unsigned Integer Elements (32-Bit) .....	6-78
6-51	vmhaddshs—Multiply-High and Add Eight Signed Integer Elements (16-Bit) .....	6-79
6-52	vmhraddshs—Multiply-High Round and Add Eight Signed Integer Elements (16-Bit) .....	6-80
6-53	vminfp—Minimum of Four Floating-Point Elements (32-Bit) .....	6-81
6-54	vminsb—Minimum of Sixteen Signed Integer Elements (8-Bit) .....	6-82
6-55	vminsh—Minimum of Eight Signed Integer Elements (16-Bit) .....	6-83
6-56	vminsw—Minimum of Four Signed Integer Elements (32-Bit) .....	6-84
6-57	vminub—Minimum of Sixteen Unsigned Integer Elements (8-Bit) .....	6-85
6-58	vminuh—Minimum of Eight Unsigned Integer Elements (16-Bit) .....	6-86
6-59	vminuw—Minimum of Four Unsigned Integer Elements (32-Bit) .....	6-87
6-60	vmladduhm—Multiply-Add of Eight Integer Elements (16-Bit) .....	6-88
6-61	vmrghb—Merge Eight High-Order Elements (8-Bit) .....	6-89
6-62	vmrghh—Merge Four High-Order Elements (16-Bit) .....	6-90
6-63	vmrghw—Merge Four High-Order Elements (32-Bit) .....	6-91
6-64	vmrglb—Merge Eight Low-Order Elements (8-Bit) .....	6-92
6-65	vmrglh—Merge Four Low-Order Elements (16-Bit) .....	6-93
6-66	vmrglw—Merge Four Low-Order Elements (32-Bit) .....	6-94
6-67	vmsummbm—Multiply-Sum of Integer Elements (8-Bit to 32-Bit) .....	6-95
6-68	vmsumshm—Multiply-Sum of Signed Integer Elements (16-Bit to 32-Bit) .....	6-96
6-69	vmsumshs—Multiply-Sum of Signed Integer Elements (16-Bit to 32-Bit) .....	6-97
6-70	vmsumubm—Multiply-Sum of Unsigned Integer Elements (8-Bit to 32-Bit) .....	6-98
6-71	vmsumuhm—Multiply-Sum of Unsigned Integer Elements (16-Bit to 32-Bit) .....	6-99
6-72	vmsumuhs—Multiply-Sum of Unsigned Integer Elements (16-Bit to 32-Bit) .....	6-100
6-73	vmulesb—Even Multiply of Eight Signed Integer Elements (8-Bit) .....	6-101
6-74	vmulesb—Even Multiply of Four Signed Integer Elements (16-Bit) .....	6-102
6-75	vmuleub—Even Multiply of Eight Unsigned Integer Elements (8-Bit) .....	6-103

## Figures

Figure Number	Title	Page Number
6-76	vmuleuh—Even Multiply of Four Unsigned Integer Elements (16-Bit) .....	6-104
6-77	vmulosb—Odd Multiply of Eight Signed Integer Elements (8-Bit) .....	6-105
6-78	vmuleuh—Odd Multiply of Four Unsigned Integer Elements (16-Bit).....	6-106
6-79	vmuloub—Odd Multiply of Eight Unsigned Integer Elements (8-Bit) .....	6-107
6-80	vmulouh—Odd Multiply of Four Unsigned Integer Elements (16-Bit) .....	6-108
6-81	vnmsubfp—Negative Multiply-Subtract of Four Floating-Point Elements (32-Bit).....	6-109
6-82	vnor—Bitwise NOR of 128-bit Vector.....	6-110
6-83	vor—Bitwise OR of 128-bit Vector.....	6-111
6-84	vperm—Concatenate Sixteen Integer Elements (8-Bit).....	6-112
6-85	How a Word is Packed to a Half Word.....	6-113
6-86	vpxpx—Pack Eight Elements (32-Bit) to Eight Elements (16-Bit).....	6-114
6-87	vpxshss—Pack Sixteen Signed Integer Elements (16-Bit) to Sixteen Signed Integer Elements (8-Bit) .....	6-115
6-88	vpxshus—Pack Sixteen Signed Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit) .....	6-116
6-89	vpxswss—Pack Eight Signed Integer Elements (32-Bit) to Eight Signed Integer Elements (16-Bit) .....	6-117
6-90	vpxswus—Pack Eight Signed Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit) .....	6-118
6-91	vpxuhum—Pack Sixteen Unsigned Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit) .....	6-119
6-92	vpxuhus—Pack Sixteen Unsigned Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit) .....	6-120
6-93	vpxuwum—Pack Eight Unsigned Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit).....	6-121
6-94	vpxuwum—Pack Eight Unsigned Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit).....	6-122
6-95	vrerp—Reciprocal Estimate of Four Floating-Point Elements (32-Bit) .....	6-124
6-96	vrerm—Round to Minus Infinity of Four Floating-Point Integer Elements (32-Bit).....	6-125
6-97	vrerfn—Nearest Round to Nearest of Four Floating-Point Integer Elements (32-Bit).....	6-126
6-98	vrerfp—Round to Plus Infinity of Four Floating-Point Integer Elements (32-Bit).....	6-127
6-99	vrerfiz—Round-to-Zero of Four Floating-Point Integer Elements (32-Bit) .....	6-128
6-100	vrlb—Left Rotate of Sixteen Integer Elements (8-Bit).....	6-129
6-101	vrlh—Left Rotate of Eight Integer Elements (16-Bit) .....	6-130
6-102	vrlw—Left Rotate of Four Integer Elements (32-Bit) .....	6-131
6-103	vrqrtefp—Reciprocal Square Root Estimate of Four Floating-Point Elements (32-Bit) .....	6-132

# Figures

Figure Number	Title	Page Number
6-104	vsel—Bitwise Conditional Select of Vector Contents(128-bit) .....	6-133
6-105	vsl—Shift Bits Left in Vector (128-Bit).....	6-134
6-106	vslb—Shift Bits Left in Sixteen Integer Elements (8-Bit).....	6-135
6-107	vsldoi—Shift Left by Bytes Specified .....	6-136
6-108	vslh—Shift Bits Left in Eight Integer Elements (16-Bit) .....	6-137
6-109	vslo—Left Byte Shift of Vector (128-Bit).....	6-138
6-110	vslw—Shift Bits Left in Four Integer Elements (32-Bit).....	6-139
6-111	vspltb—Copy Contents to Sixteen Elements (8-Bit) .....	6-140
6-112	vsplth—Copy Contents to Eight Elements (16-Bit).....	6-141
6-113	vspltisb—Copy Value into Sixteen Signed Integer Elements (8-Bit) .....	6-142
6-114	vspltish—Copy Value to Eight Signed Integer Elements (16-Bit).....	6-143
6-115	vspltisw—Copy Value to Four Signed Elements (32-Bit) .....	6-144
6-116	vspltw—Copy contents to Four Elements (32-Bit).....	6-145
6-117	vsr—Shift Bits Right for Vectors (128-Bit) .....	6-147
6-118	vsrab—Shift Bits Right in Sixteen Integer Elements (8-Bit).....	6-148
6-119	vsrah—Shift Bits Right for Eight Integer Elements (16-Bit).....	6-149
6-120	vsraw—Shift Bits Right in Four Integer Elements (32-Bit) .....	6-150
6-121	vsrb—Shift Bits Right in Sixteen Integer Elements (8-Bit).....	6-151
6-122	vsrh—Shift Bits Right for Eight Integer Elements (16-Bit) .....	6-152
6-123	vsro—Vector Shift Right Octet .....	6-153
6-124	vsrw—Shift Bits Right in Four Integer Elements (32-Bit) .....	6-154
6-125	vsubcuw—Subtract Carryout of Four Unsigned Integer Elements (32-Bit).....	6-155
6-126	vsubfp—Subtract Four Floating Point Elements (32-Bit) .....	6-156
6-127	vsubsb—Subtract Sixteen Signed Integer Elements (8-Bit).....	6-157
6-128	vsubsh—Subtract Eight Signed Integer Elements (16-Bit) .....	6-158
6-129	vsubsw—Subtract Four Signed Integer Elements (32-Bit) .....	6-159
6-130	vsububm—Subtract Sixteen Integer Elements (8-Bit).....	6-160
6-131	vsubub—Subtract Sixteen Unsigned Integer Elements (8-Bit) .....	6-161
6-132	vsubuhm—Subtract Eight Integer Elements (16-Bit) .....	6-162
6-133	vsubuh—Subtract Eight Signed Integer Elements (16-Bit).....	6-163
6-134	vsubuw—Subtract Four Integer Elements (32-Bit) .....	6-164
6-135	vsubuw—Subtract Four Signed Integer Elements (32-Bit).....	6-165
6-136	vsumsw—Sum Four Signed Integer Elements (32-Bit) .....	6-166
6-137	vsum2sw—Two Sums in the Four Signed Integer Elements (32-Bit).....	6-167
6-138	vsum4sb—Four Sums in the Integer Elements (32-Bit) .....	6-168
6-139	vsum4sh—Four Sums in the Integer Elements (32-Bit) .....	6-169
6-140	vsum4ub—Four Sums in the Integer Elements (32-Bit) .....	6-170
6-141	vupkhp—Unpack High-Order Elements (16 bit) to Elements (32-Bit).....	6-171
6-142	vupkhsb—Unpack High-Order Signed Integer Elements (8-Bit) to Signed Integer Elements (16-Bit) .....	6-172

## Figures

<b>Figure Number</b>	<b>Title</b>	<b>Page Number</b>
6-143	vupkhsh—Unpack Signed Integer Elements (16-Bit) to Signed Integer Elements (32-Bit) .....	6-173
6-144	vupklsx—Unpack Low-order Elements (16-Bit) to Elements (32-Bit) .....	6-174
6-145	vupklsb—Unpack Low-Order Elements (8-Bit) to Elements (16-Bit) .....	6-175
6-146	vupklsh—Unpack Low-Order Signed Integer Elements (16-Bit) to Signed Integer Elements (32-Bit) .....	6-176
6-147	vxor—Bitwise XOR (128-Bit) .....	6-177



## Tables

Table Number	Title	Page Number
i	Acronyms and Abbreviated Terms .....	xxiv
ii	Terminology Conventions .....	xxvii
iii	Instruction Field Conventions .....	xxvii
2-1	VSCR Field Descriptions .....	2-5
2-2	VRSAVE Bit Settings .....	2-7
2-3	CR6 Field's Bit Settings for Vector Compare Instructions .....	2-8
2-4	MSR Bit Settings .....	2-10
3-1	Memory Operand Alignment .....	3-2
3-2	Effective Address Modifications .....	3-5
4-1	Vector Integer Arithmetic Instructions .....	4-6
4-2	CR6 Field Bit Settings for Vector Integer Compare Instructions .....	4-13
4-3	Vector Integer Compare Instructions .....	4-14
4-4	Vector Integer Logical Instructions .....	4-16
4-5	Vector Integer Rotate Instructions .....	4-16
4-6	Vector Integer Shift Instructions .....	4-17
4-7	Floating-Point Arithmetic Instructions .....	4-19
4-8	Floating-Point Multiply-Add Instructions .....	4-21
4-9	Floating-Point Rounding and Conversion Instructions .....	4-21
4-10	Common Mathematical Predicates .....	4-23
4-11	Other Useful Predicates .....	4-23
4-12	Floating-Point Compare Instructions .....	4-24
4-13	Floating-Point Estimate Instructions .....	4-25
4-14	Effective Address Alignment .....	4-26
4-15	Integer Load Instructions .....	4-28
4-16	Vector Load Instructions Supporting Alignment .....	4-29
4-17	Shift Values for lvsl Instruction .....	4-29
4-18	Shift Values for lvsr Instruction .....	4-29
4-19	Integer Store Instructions .....	4-30
4-20	Vector Pack Instructions .....	4-32
4-21	Vector Unpack Instructions .....	4-34
4-22	Vector Merge Instructions .....	4-35
4-23	Vector Splat Instructions .....	4-36
4-24	Vector Permute Instruction .....	4-36
4-25	Vector Select Instruction .....	4-37
4-26	Vector Shift Instructions .....	4-37

## Tables

Table Number	Title	Page Number
4-27	Coding Various Shifts and Rotates with the vsidoi Instruction.....	4-38
4-28	Move to/from Condition Register Instructions .....	4-40
4-29	Simplified Mnemonics for Data Stream Touch ( <b>dst</b> ).....	4-40
4-30	User-Level Cache Instructions .....	4-42
5-1	AltiVec Unavailable Exception—Register Settings.....	5-11
5-2	Exception Priorities (Synchronous/Precise Exceptions).....	5-12
6-1	Instruction Syntax Conventions .....	6-2
6-2	Notation and Conventions.....	6-2
6-3	Instruction Field Conventions .....	6-7
6-4	Precedence Rules .....	6-7
6-5	Special Values of the Element in vB .....	6-67
6-6	Special Values of the Element in vB .....	6-69
6-7	Special Values of the Element in vB .....	6-123
6-8	Special Values of the Element in vB .....	6-132
A-1	Instruction Sorted by Mnemonic in Decimal Format .....	A-1
B-1	Instructions Sorted by Mnemonic in Binary Format.....	B-1
C-1	Instructions Sorted by Opcode in Decimal Format.....	C-1
D-1	Instructions Sorted by Opcode in Binary Format .....	D-1
E-1	VA-Form.....	E-1
E-2	VX-Form.....	E-2
E-3	X-Form.....	E-5
E-4	VXR-Form .....	E-6
F-1	AltiVec Instruction Set Legend .....	F-1

## About This Book

---

The primary objective of this manual is to help programmers provide software that is compatible with processors that implement the PowerPC architecture and the AltiVec™ technology. This book describes how the AltiVec technology relates to the 32-bit portions of the PowerPC architecture.

To locate any published errata or updates for this document, refer to the web at <http://www.motorola.com/semiconductors>.

This book is one of two that discuss the AltiVec technology. The two books are as follows.

- *AltiVec Technology Programming Interface Manual (AltiVec PIM)* is a reference guide for high-level programmers. The AltiVec PIM describes how programmers can access AltiVec functionality from programming languages such as C and C++. The AltiVec PIM defines a programming model for use with the AltiVec instruction set. Processor that implement the PowerPC architecture use the AltiVec instruction set as an extension of the PowerPC instruction set.
- *AltiVec Technology Programming Environments Manual (AltiVec PEM)* is used as a reference guide for assembler programmers. The AltiVec PEM uses a standardized format instruction to describe each instruction, showing syntax, instruction format, register translation language (RTL) code that describes how the instruction works, and a listing of which, if any, registers are affected. At the bottom of each instruction entry is a figure that shows the operations on elements within source operands and where the results of those operations are placed in the destination operand.

Because it is important to distinguish between the levels of the PowerPC architecture to ensure compatibility across multiple platforms, those distinctions are shown clearly throughout this book. This document stays consistent with the PowerPC architecture in referring to three levels, or programming environments, which are as follows:

- • PowerPC user instruction set architecture (UISA)—The UISA defines the level of the architecture to which user-level software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, memory conventions, and the memory and programming models seen by application programmers.
- ▼ • PowerPC virtual environment architecture (VEA)—The VEA, which is the smallest component of the PowerPC architecture, defines additional user-level functionality that falls outside typical user-level software requirements. The VEA describes the memory model for an environment in which multiple processors or other devices can

access external memory and defines aspects of the cache model and cache control instructions from a user-level perspective. VEA resources are particularly useful for optimizing memory accesses and for managing resources in an environment in which other processors and other devices can access external memory.

Implementations that conform to the VEA also conform to the UISA but may not necessarily adhere to the OEA.

- PowerPC operating environment architecture (OEA)—The OEA defines supervisor-level resources typically required by an operating system. It defines the memory management model, supervisor-level registers, and the exception model. Implementations that conform to the OEA also conform to the UISA and VEA.

Most of the discussions on the AltiVec technology are at the UISA level. The level of the architecture to which text refers is indicated in the outer margin, using the conventions shown in Section , “Conventions,” on page -xxiii.

For ease in reference, this book and the processor user’s manuals have arranged the architecture information into topics that build upon one another, beginning with a description and complete summary of registers and instructions (for all three environments) and progressing to more specialized topics such as the cache, exception, and memory management models. As such, chapters may include information from multiple levels of the architecture, but when discussing OEA and VEA, the level is noted in the text.

It is beyond the scope of this manual to describe individual AltiVec technology implementations on processors that implement the PowerPC architecture. It must be kept in mind that each processor that implements the PowerPC architecture and AltiVec technology is unique in its implementation.

The information in this book is subject to change without notice, as described in the disclaimers on the title page of this book. As with any technical documentation, it is the readers’ responsibility to be sure they are using the most recent version of the documentation. For more information, contact your sales representative or visit our web site at <http://www.mot.com/semiconductors>.

## Audience

This manual is intended for system software and hardware developers and application programmers who want to develop products using the AltiVec technology extension to the PowerPC architecture. It is assumed that the reader understands operating systems, microprocessor system design, and the basic principles of RISC processing and details of the PowerPC architecture.

This book describes how the AltiVec technology interacts with the 32-bit portions of the PowerPC architecture

## Organization

Following is a summary and a brief description of the major sections of this manual:

- Chapter 1, “Overview,” is useful for those who want a general understanding of the features and functions of the AltiVec technology. This chapter provides an overview of how the AltiVec technology defines the register set, operand conventions, addressing modes, instruction set, cache model, and exception model.
- Chapter 2, “AltiVec Register Set,” is useful for software engineers who need to understand the PowerPC programming model for the three programming environments. The chapter also discusses the functionality of the AltiVec technology registers and how they interact with the other PowerPC registers.
- Chapter 3, “Operand Conventions,” describes how the AltiVec technology interacts with the PowerPC conventions for storing data in memory, including information regarding alignment, single-precision floating-point conventions, and big- and little-endian byte ordering.
- Chapter 4, “Addressing Modes and Instruction Set Summary,” provides an overview of the AltiVec technology addressing modes and a brief description of the AltiVec technology instructions organized by function.
- Chapter 5, “Cache, Exceptions, and Memory Management,” provides a discussion of the cache and memory model defined by the VEA and aspects of the cache model that are defined by the OEA. It also describes the exception model defined in the UISA.
- Chapter 6, “AltiVec Instructions,” functions as a handbook for the AltiVec instruction set. Instructions are sorted by mnemonic. Each instruction description includes the instruction formats and figures where it helps in understanding what the instruction does.
- Appendices A, B, C, D, E, F, and G list all of the AltiVec instructions, grouped according to mnemonic, opcode, and form, in both decimal and binary order.
- Appendix G, “User’s Manual Revision History,” describes changes since the previous revision of this document.
- This manual also includes a glossary and an index.

## Suggested Reading

This section lists additional reading that provides background for the information in this manual as well as general information about the AltiVec technology and PowerPC architecture.

## General Information

The following documentation, available through Morgan-Kaufmann Publishers, 340 Pine Street, Sixth Floor, San Francisco, CA, provides useful information about the PowerPC architecture and computer architecture in general:

- *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.  
For updates to the specification, see <http://www.austin.ibm.com/tech/ppc-chg.html>.
- *PowerPC Microprocessor Common Hardware Reference Platform: A System Architecture*, by Apple Computer, Inc., International Business Machines, Inc., and Motorola, Inc.
- *Computer Architecture: A Quantitative Approach*, Second Edition, by John L. Hennessy and David A. Patterson
- *Computer Organization and Design: The Hardware/Software Interface*, Second Edition, David A. Patterson and John L. Hennessy

## Related Documentation

Motorola documentation is available from the sources listed on the back cover of this manual; the document order numbers are included in parentheses for ease in ordering:

- *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture* (Programming Environments Manual)—Describes resources defined by the PowerPC architecture (documentation order number: MPCFP32B/AD).
- User's manuals—These books provide details about individual implementations and are intended for use with the *Programming Environments Manual*.
- Addenda/errata to user's manuals—Because some processors have follow-on parts an addendum is provided that describes the additional features and functionality changes. These addenda are intended for use with the corresponding user's manuals.
- Hardware specifications—Hardware specifications provide specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations.
- Technical summaries—Each device has a technical summary that provides an overview of its features. This document is roughly the equivalent to the overview (Chapter 1) of an implementation's user's manual.
- Application notes—These short documents address specific design issues useful to programmers and engineers working with Motorola processors.

Additional literature is published as new processors become available. For a current list of documentation, refer to <http://www.motorola.com/semiconductors>.

## Conventions

This document uses the following notational conventions:

cleared/set	When a bit takes the value zero, it is said to be cleared; when it takes a value of one, it is said to be set.
<b>mnemonics</b>	Instruction mnemonics are shown in lowercase bold.
<i>italics</i>	Italics indicate variable command parameters, for example, <b>bcctrx</b> . Book titles in text are set in italics
0x0	Prefix to denote hexadecimal number
0b0	Prefix to denote binary number
<b>rA, rB</b>	Instruction syntax used to identify a source general-purpose register (GPR)
<b>rD</b>	Instruction syntax used to identify a destination GPR
<b>frA, frB, frC</b>	Instruction syntax used to identify a source floating-point register (FPR)
<b>frD</b>	Instruction syntax used to identify a destination FPR
REG[FIELD]	Abbreviations for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register.
<b>vA, vB, vC</b>	Instruction syntax used to identify a source vector register (VR)
<b>vD</b>	Instruction syntax used to identify a destination VR
<i>x</i>	In some contexts, such as signal encodings, an unitalicized <i>x</i> indicates a don't care.
<i>x</i>	An italicized <i>x</i> indicates an alphanumeric variable.
<i>n</i>	An italicized <i>n</i> indicates an numeric variable.
¬	NOT logical operator
&	AND logical operator
	OR logical operator
<b>U</b>	This symbol identifies text that is relevant with respect to the PowerPC user instruction set architecture (UISA). This symbol is used both for information that can be found in the UISA specification as well as for explanatory information related to that programming environment.




This symbol identifies text that is relevant with respect to the PowerPC virtual environment architecture (VEA). This symbol is used both for information that can be found in the VEA specification as well as for explanatory information related to that programming environment.



This symbol identifies text that is relevant with respect to the PowerPC operating environment architecture (OEA). This symbol is used both for information that can be found in the OEA specification as well as for explanatory information related to that programming environment.



Indicates functionality defined by the AltiVec technology.



Indicates reserved bits or bit fields in a register. Although these bits may be written to as ones or zeros, they are always read as zeros.

Additional conventions used with instruction encodings are described in Section 6.1, “Instruction Formats.”

## Acronyms and Abbreviations

Table i contains acronyms and abbreviations that are used in this document. Note that the meanings for some acronyms (such as SDR1 and XER) are historical, and the words for which an acronym stands may not be intuitively obvious.

**Table i. Acronyms and Abbreviated Terms**

Term	Meaning
AltiVec PEM	<i>AltiVec Technology Programming Environments Manual</i>
AltiVec PIM	<i>AltiVec Technology Programming Interface Manual</i>
ALU	Arithmetic logic unit
BAT	Block address translation
CR	Condition register
CTR	Count register
DABR	Data address breakpoint register
DAR	Data address register
DBAT	Data BAT
DEC	Decrementer register
DSISR	Register used for determining the source of a DSI exception
EA	Effective address
ECC	Error checking and correction



## Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
FPR	Floating-point register
FPSCR	Floating-point status and control register
FPU	Floating-point unit
GPR	General-purpose register
IABR	Instruction address breakpoint register
IBAT	Instruction BAT
IEEE	Institute of Electrical and Electronics Engineers
ITLB	Instruction translation lookaside buffer
IU	Integer unit
L2	Secondary cache
L3	Level 3 cache
LIFO	Last-in-first-out
LR	Link register
LRU	Least recently used
LSB	Least-significant byte
lsb	Least-significant bit
LSU	Load/store unit
LSQ	Least-significant quad-word
lsq	Least-significant quad-word
MESI	Modified/exclusive/shared/invalid—cache coherency protocol
MMCR <sub>n</sub>	Monitor mode control registers
MMU	Memory management unit
MSB	Most-significant byte
msb	Most-significant bit
MSQ	Most-significant quad-word
msq	Most-significant quad-word
MSR	Machine state register
NaN	Not a number
NIA	Next instruction address
No-op	No operation
OEA	Operating environment architecture
PEM	<i>Programming Environments Manual For 32-Bit Implementations of the PowerPC Architecture</i>
PMC <sub>n</sub>	Performance monitor counter registers
PTE	Page table entry

## Table i. Acronyms and Abbreviated Terms (continued)

Term	Meaning
PTEG	Page table entry group
PVR	Processor version register
RISC	Reduced instruction set computing
RTL	Register transfer language
RWITM	Read with intent to modify
RWNITM	Read with no intent to modify
SDA	Sampled data address register
SDR1	Register that specifies the page table base address for virtual-to-physical address translation
SIA	Sampled instruction address register
SIMM	Signed immediate value
SPR	Special-purpose register
SR $n$	Segment register
SRR0	Machine status save/restore register 0
SRR1	Machine status save/restore register 1
STE	Segment table entry
TB	Time base facility
TBL	Time base lower register
TBU	Time base upper register
TLB	Translation lookaside buffer
UIMM	Unsigned immediate value
UISA	User instruction set architecture
UMMCR $n$	User monitor mode control registers
UPMC $n$	User performance monitor counter registers
VA	Virtual address
VEA	Virtual environment architecture
VPU	Vector permute unit
VR	Vector register
VSCR	Vector status and control register
VTQ	Vector touch queue
XER	Register used for indicating conditions such as carries and overflows for integer operations

## Terminology Conventions

Table ii lists certain terms used in this manual that differ from the architecture terminology conventions.

**Table ii. Terminology Conventions**

The Architecture Specification	This Manual
Data storage interrupt (DSI)	DSI exception
Extended mnemonics	Simplified mnemonics
Fixed-point unit (FXU)	Integer unit (IU)
Instruction storage interrupt (ISI)	ISI exception
Interrupt	Exception
Privileged mode (or privileged state)	Supervisor-level privilege
Problem mode (or problem state)	User-level privilege
Real address	Physical address
Relocation	Translation
Storage (locations)	Memory
Storage (the act of)	Access
Store in	Write back
Store through	Write through

Table iii describes instruction field notation conventions used in this manual.

**Table iii. Instruction Field Conventions**

The Architecture Specification	Equivalent to:
BA, BB, BT	<b>crbA, crbB, crbD</b> (respectively)
BF, BFA	<b>crfD, crfS</b> (respectively)
D	d
DS	ds
FLM	FM
FRA, FRB, FRC, FRT, FRS	<b>frA, frB, frC, frD, frS</b> (respectively)
FXM	CRM
RA, RB, RT, RS	<b>rA, rB, rD, rS</b> (respectively)
SI	SIMM
U	IMM
UI	UIMM
VA, VB, VT, VS	<b>vA, vB, vD, vS</b> (respectively)

**Table iii. Instruction Field Conventions (continued)**

The Architecture Specification	Equivalent to:
VEC	AltiVec technology
I, II, III	0...0 (shaded)

# Chapter 1

## Overview

This chapter provides an overview of AltiVec™ technology, including general concepts which help in understanding the features that AltiVec technology provides. There is also information on how AltiVec technology works with PowerPC architecture.

### 1.1 Overview

AltiVec™ technology provides a software model that accelerates the performance of various software applications as it runs on reduced instruction set computing (RISC) microprocessors. AltiVec technology extends the instruction set architecture (ISA) of PowerPC architecture. AltiVec ISA is based on separate vector/SIMD-style (single instruction stream, multiple data streams) execution units that have high data parallelism. That is, AltiVec technology operates on multiple data items in a single instruction which allows for a highly efficient way to process large quantities of information. High degrees of parallelism are achievable with simple in-order instruction dispatch and low-instruction time processing. However, the ISA is designed so as not to impede additional parallelism through dispatch to multiple execution units or multithreaded execution unit pipelines.

AltiVec technology is an architecture that defines a set of registers and execution units which can be used in conjunction with the PowerPC architecture. All instructions are designed to be easily pipelined with pipeline latencies no greater than the scalar, double-precision, floating-point multiply-add. There are no operating mode switches which make interleaving of instructions with the existing floating-point and integer instructions possible. The vector unit minimizes exceptions and has few shared resources. This requires it to be tightly synchronized with other execution units that prevent delays in executing instructions.

AltiVec technology's SIMD-style extension provides an approach to accelerating the processing of data streams. That is, in SIMD parallel processing, the vector unit will fetch and interpret instructions and process multiple pieces of data simultaneously. By processing whole streams of data at once, it provides a fast and efficient way to manipulate large quantities of information. AltiVec instructions provide a significant speedup for communications, multimedia, and other performance-driven applications by using the data-level parallelism and keeping processing of data to the vector register file. By having separate register files, the execution units data accesses by different register files can be

done concurrently. The data stream engine in AltiVec supports data-intensive prefetching, minimizing latency in memory access bottlenecks. By using the SIMD parallelism in AltiVec technology, performance can be accelerated on processors that implement the PowerPC architecture to a level that allows real-time processing of one or more data streams at the same time.

A majority of audio and visual applications require no more than 8- or 16-bit data types to represent satisfactory color and sound. AltiVec ISA can help accelerate the processing of the following types of applications:

- Voice over IP (VoIP). VoIP transmits voice as compressed digital data packets over the Internet.
- Access Concentrators/DSLAMS. An access concentrator strips data traffic off POTS lines and inserts it onto the Internet. Digital subscriber loop access multiplexer (DSLAM) pulls data off at a switch and immediately routes it to the Internet. This allows it to concentrate ADSL digital traffic at the switch and off-load the network.
- Speech recognition. Speech processing allows voice recognition for use in applications such as directory assistance and automatic dialing.
- Voice/sound processing (audio encode and decode): Voice processing uses signal processing to improve sound quality on lines.
- Communications:
  - Multi-channel modems
  - Modem banks can use AltiVec technology to replace signal processors in DSP farms.
- 2D and 3D graphics: arcade-type games
- Image and video processing: JPEG, filters
- Echo cancellation. Echo cancellation is used to eliminate echo on long delay calls (250–500 milliseconds, as in satellite communications).
- Array number processing
- Basestation Processing: Cellular basestation compresses digital voice data for transmission within the Internet.
- Video conferencing: H.261, H.263

In this document, the term ‘implementation’ refers to a hardware device (typically a microprocessor) that complies with PowerPC architecture.

AltiVec technology can be used as an extension to various RISC microprocessors; however, in this book it is discussed within the context of PowerPC architecture, described as follows:

- Programming model
  - Instruction set. The AltiVec instruction set specifies instructions that extend the PowerPC instruction set. These instructions are organized similar to PowerPC instructions (vector integer, vector floating-point, vector load/store, and vector permutation and formatting instructions). The specific instructions, and the forms used for encoding them, are provided in Appendix A, “Instruction Set.”
  - Register set. The AltiVec programming model defines new AltiVec registers, additions to the PowerPC register set, and how existing PowerPC registers are affected by the AltiVec technology. The model also addresses memory conventions including details regarding the byte ordering for quad words.
- Memory model. AltiVec technology specifies additional cache management instructions. That is, AltiVec instructions can control software-directed data prefetching.
- Exception model. AltiVec technology provides very few exceptions, so processing is efficient. Among the few exceptions are an AltiVec unavailable (VUI) exception and a DSI exception.
- Memory management model. The memory model for AltiVec technology is the same as for PowerPC architecture. AltiVec memory accesses are always assumed to be aligned. If an operand is misaligned, additional AltiVec instructions can be used to ensure that the operand is placed correctly in the vector register.
- Time-keeping model. The PowerPC time-keeping model is not affected by AltiVec technology.

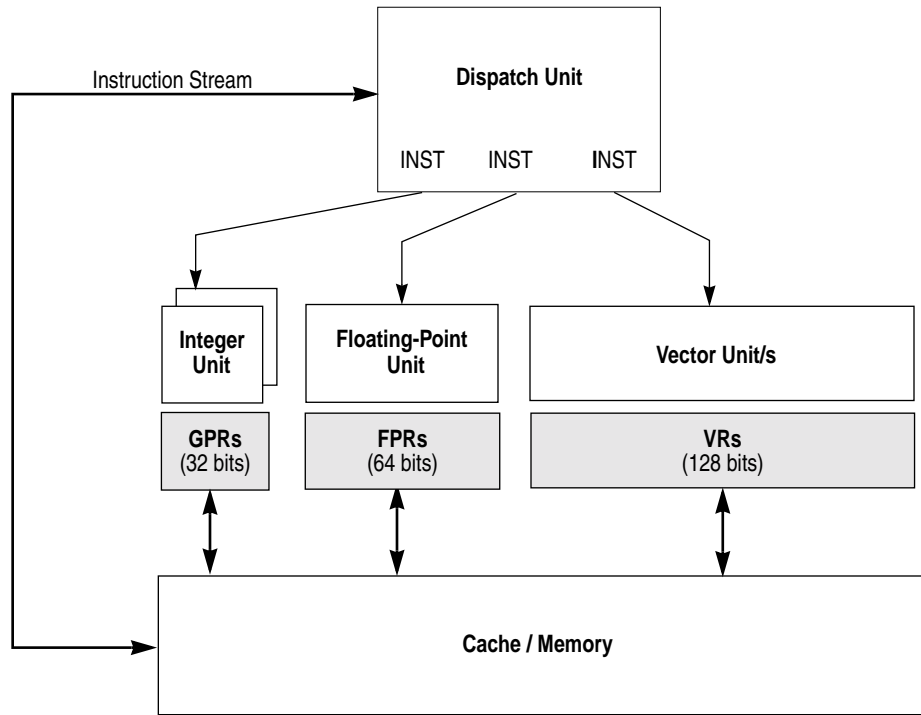
To locate published errata or updates for this document, refer to the website at <http://www.motorola.com/semiconductors>.

## 1.2 AltiVec Technology Overview

AltiVec technology expands PowerPC architecture through the addition of a 128-bit vector execution unit, which operates concurrently with the existing integer- and floating-point units. The dispatch unit can issue more than one instruction at a time so there is no penalty for mingling different types of instructions. A new vector execution unit can provide both a vector permute unit (VPERM) and vector arithmetic logical unit (VALU). By having a separate permute unit, data reorganization instructions can proceed concurrently with arithmetic instructions.

AltiVec technology can be thought of as a set of registers and execution units that can be added to PowerPC architecture in a manner analogous to the addition of floating-point units. Floating-point units were added to provide support for high-precision scientific calculations, and AltiVec technology is added to PowerPC architecture to accelerate the next level of performance-driven, high-bandwidth communications and computing

applications. Figure 1-1 provides a high-level overview of the PowerPC architecture with the AltiVec technology.



**Figure 1-1. Overview of PowerPC architecture with AltiVec Technology**

AltiVec technology is purposefully simple so that there are minimal exceptions, no hardware misaligned access support, and no complex functions. AltiVec technology is scaled down to the necessary pieces only, in order to facilitate efficient cycle time, latency, and throughput on hardware implementations.

AltiVec technology defines the following:




- Fixed 128-bit-wide vector length that can be subdivided into sixteen 8-bit bytes, eight 16-bit half words, or four 32-bit words
- Vector register file (VRF) architecturally separate from floating-point registers (FPRs) and general-purpose registers (GPRs)
- Vector integer and floating-point arithmetic
- Four operands for most instructions (three source operands and one result)
- Saturation clamping (that is, unsigned results are clamped to zero on underflow and to the maximum positive integer value ( $2^n-1$ , for example, 255 for byte fields) on overflow. For signed results, saturation clamps results to the smallest representable negative number ( $-2^{n-1}$ , for example, -128 for byte fields) on underflow, and to the largest representable positive number ( $2^{n-1}-1$ , for example, +127 for byte fields) on overflow)



- Operations selected based on utility to digital signal processing algorithms (including 3D).
- AltiVec instructions provide a vector compare and select mechanism to implement conditional execution as the preferred way to control data flow in AltiVec programs.
- Instructions that enhance the cache/memory interface

## 1.2.1 Levels of AltiVec ISA

AltiVec ISA follows the layering of PowerPC architecture. PowerPC architecture has three levels, defined as follows:

- User instruction set architecture (UISA) —The UISA defines the level of the architecture to which user-level (referred to as problem state in the architecture specification) software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, floating-point memory conventions, and exception model as seen by user programs, and the memory and programming models. The icon shown in the margin identifies text that is relevant to the UISA. 
- Virtual environment architecture (VEA)—The VEA defines additional user-level functionality that falls outside typical user-level software requirements. The VEA describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time base facility from a user-level perspective. The icon shown in the margin identifies text that is relevant to the VEA.   
Implementations that conform to the VEA also adhere to the UISA, but may not necessarily adhere to the OEA.
- Operating environment architecture (OEA)—The OEA defines supervisor-level (referred to as privileged state in the architecture specification) resources typically required by an operating system. The OEA defines the memory management model, supervisor-level registers, synchronization requirements, and the exception model. The OEA also defines the time base feature from a supervisor-level perspective. The icon shown in the margin identifies text that is relevant to the OEA.   
Implementations that conform to the OEA also conform to the UISA and VEA.

AltiVec technology defines instructions at the UISA and VEA levels. There are no AltiVec instructions defined at the OEA level. The distinctions between the levels are noted in the text throughout the document. This book describes the 32-bit PowerPC architecture mode, and instructions are described from a 32-bit perspective.

## 1.2.2 Features Not Defined by AltiVec ISA

Because flexibility is an important design goal of AltiVec technology, there are many aspects of the microprocessor design, typically relating to the hardware implementation, that AltiVec ISA does not define. For example, the number and the nature of execution units are not defined. AltiVec ISA is a vector/SIMD architecture, and as such makes it easier to implement pipelining instructions and parallel execution units to maximize instruction throughput. However, AltiVec ISA does not define the internal hardware details of implementations. For example, one processor may use a simple implementation having two vector execution units, whereas another may provide a bigger, faster microprocessor design with several concurrently pipelined vector arithmetic logical units (ALUs) with separate load/store units (LSUs) and prefetch units.


## 1.3 AltiVec Architectural Model

This section provides overviews of aspects defined by AltiVec ISA, following the same order as the rest of this book. The topics are as follows:

- Registers and programming model
- Operand conventions
- Addressing modes and instruction set
- Cache, exceptions, and memory management models

### 1.3.1 AltiVec Registers and Programming Model

In AltiVec technology, the ALU operates on from one to three source vectors and produces a single destination vector on each instruction. The ALU is a SIMD-style arithmetic unit that performs the same operation on all the data elements comprising each vector. This scheme allows efficient code scheduling in a highly parallel processor. Load and store instructions are the only instructions that transfer data between registers and memory. The vector unit and vector register file are shown in Figure 1-2.



**Figure 1-2. AltiVec Top-Level Diagram**

The vector unit is a SIMD-style unit in which an instruction performs operations in parallel with the data elements that comprise each vector. Architecturally, the vector register file (VRF) is separate from the GPRs and FPRs. The AltiVec programming model incorporates the 32 registers of the VRFs; each register is 128 bits wide.


## 1.3.2 Operand Conventions

Operand conventions define how data is stored in vector registers and memory.

### 1.3.2.1 Byte Ordering

The default mapping for AltiVec ISA is PowerPC big-endian, but AltiVec ISA provides the option of operating in either big- or little-endian mode. The endian support of PowerPC architecture does not address any data element larger than a double word; the basic memory unit for vectors is a quad word.


Big-endian byte ordering is shown in Figure 1-3



**Figure 1-3. Big-Endian Byte Ordering for a Vector Register**

As shown in Figure 1-3, the elements in vector registers are numbered using big-endian byte ordering. For example, the high-order (or most significant) byte element is numbered 0 and the low-order (or least significant) byte element is numbered 15.

When defining high order and low order for elements in a vector register, be careful not to confuse its meaning based on the bit numbering. That is, in Figure 1-4, the high-order half word for word 0 (bits 0–31) would be half word 0 (bits 0–15), and the low-order half word for word 0 would be half word 1 (bits 16–31).



**Figure 1-4. Bit Ordering**

In big-endian mode, an AltiVec quad word load instruction for which the effective address (EA) is quad-word aligned places the byte addressed by EA into byte element 0 of the target vector register. The byte addressed by EA + 1 is placed in byte element 1, and so forth. Similarly, an AltiVec quad word store instruction for which the EA is quad word-aligned places byte element 0 of the source vector register into the byte addressed by EA. Byte element 1 is placed into the byte addressed by EA + 1, and so forth.

### 1.3.2.2 Floating-Point Conventions

AltiVec ISA basically has two modes for floating-point, that is a Java-/IEEE-/C9X-compliant mode or a possibly faster non-Java/non-IEEE mode. AltiVec ISA conforms to the Java Language Specification 1 (hereafter referred to as Java), that is a subset of the default environment specified by the IEEE standard (ANSI/IEEE Standard 754-1985, IEEE Standard for Binary Floating-Point Arithmetic). For aspects of floating-point behavior that are not defined by Java but are defined by the IEEE standard, AltiVec ISA conforms to the IEEE standard. For aspects of floating-point behavior that are defined neither by Java nor by the IEEE standard but are defined by the C9X Floating-Point

Proposal WG14/N546 X3J11/96-010 (Draft 2/26/96) (hereafter referred to as C9X), AltiVec ISA conforms to C9X when in Java-compliant mode.

### 1.3.3 AltiVec Addressing Modes

As with PowerPC instructions, AltiVec instructions are encoded as single-word (32-bit) instructions. Instruction formats are consistent among all instruction types, permitting decoding to be parallel with operand accesses. This fixed instruction length and consistent format simplifies instruction pipelining. AltiVec load, store, and stream prefetch instructions use secondary opcodes in primary opcode 31 (0b011111). AltiVec ALU-type instructions use primary opcode 4 (0b000100).

AltiVec ISA supports both intraelement and interelement operations. In an intraelement operation, elements work in parallel with the corresponding elements from multiple source operand registers and place the results in the corresponding fields in the destination operand register. An example of an intraelement operation is the Vector Add Signed Word Saturate (**vaddsws**) instruction shown in Figure 1-5




Figure 1-5. Intraelement Example, vaddsws

In this example, the sixteen elements (8 bits per element) in register vA are added to the corresponding sixteen elements (8 bits per element) in register vB and the sixteen results are placed in the corresponding elements in register vD.

In interelement operations data paths cross over. That is, different elements from each source operand are used in the resulting destination operand. An example of an interelement operation is the Vector Permute (**vperm**) instruction shown in Figure 1-6.




Figure 1-6. Interelement Example, vperm

In this example, **vperm** allows any byte in the two source vector registers (**vA** and **vB**) to be copied to any byte in the destination vector register, **vD**. The bytes in a third source vector register (**vC**) specify from which byte in the first two source vector registers the corresponding target byte is to be copied. So in the interelement example, the elements from the source vector registers do not have corresponding elements that operate on the destination register.

Most arithmetic and logical instructions are intraelement operations. The crossover data paths have been restricted as much as possible to the interelement manipulation instructions (unpack, pack, permute, etc.) with the idea to implement the ALU and shift/permute as separate execution units. The following list of instructions distinguishes between interelement and intraelement instructions:

- Vector intraelement instructions
  - Vector integer instructions
    - Vector integer arithmetic instructions
    - Vector integer compare instructions
    - Vector integer rotate and shift instructions
  - Vector floating-point instructions
    - Vector floating-point arithmetic instructions
    - Vector floating-point rounding and conversion instructions
    - Vector floating-point compare instruction
    - Vector floating-point estimate instructions
  - Vector memory access instructions
- Vector interelement instructions
  - Vector alignment support instructions
  - Vector permutation and formatting instructions
    - Vector pack instructions
    - Vector unpack instructions
    - Vector merge instructions
    - Vector splat instructions
    - Vector permute instructions
    - Vector shift left/right instructions

### 1.3.4 AltiVec Instruction Set

Although these categories are not defined by AltiVec ISA, AltiVec instructions can be grouped as follows:

- U** • Vector integer arithmetic instructions—These instructions are defined by the UISA. They include computational, logical, rotate, and shift instructions.
  - Vector integer arithmetic instructions
  - Vector integer compare instructions
  - Vector integer logical instructions
  - Vector integer rotate and shift instructions
- U** • Vector floating-point arithmetic instructions—These include floating-point arithmetic instructions defined by the UISA.
  - Vector floating-point arithmetic instructions
  - Vector floating-point multiply/add instructions
  - Vector floating-point rounding and conversion instructions
  - Vector floating-point compare instruction
  - Vector floating-point estimate instructions
- U** • Vector load and store instructions—These include load and store instructions for vector registers defined by the UISA.
- U** • Vector permutation and formatting instructions—These instructions are defined by the UISA.
  - Vector pack instructions
  - Vector unpack instructions
  - Vector merge instructions
  - Vector splat instructions
  - Vector permute instructions
  - Vector select instructions
- U** • Vector shift instructions
- Processor control instructions—These instructions are used to read and write from the AltiVec status and control register (VSCR). These instructions are defined by the UISA.
- V** • Memory control instructions—These instructions are used for managing of caches (user level and supervisor level). The instructions are defined by VEA and include data stream instructions.

### **1.3.5 AltiVec Cache Model**

AltiVec ISA defines several instructions for enhancements to cache management. These instructions allow software to indicate to the cache hardware how it should prefetch and prioritize writeback of data. The AltiVec ISA does not define hardware aspects of cache implementations.

### **1.3.6 AltiVec Exception Model**

AltiVec vector instructions generate very few exceptions. Data stream instructions will never cause an exception themselves. Vector load and store instructions that attempt to access a direct-store segment will cause a DSI exception.

The AltiVec unit does not report IEEE exceptions; there are no status flags and the unit has no architecturally visible traps. Default results are produced for all exception conditions as specified first by the Java specification. If no default exists, the IEEE standard's default is used. Then, if no default exists, the C9X default is used.

Exceptions have been minimized so that the vector unit does not have to be tightly synchronized with the existing floating-point and integer units. By simplifying the communications path with other units there can be fine grain interleaving of instructions that increases the instruction through-put.

### **1.3.7 Memory Management Model**

In a processor that implement the PowerPC architecture the MMU's primary functions are to translate logical (effective) addresses to physical addresses for memory accesses and I/O accesses (most I/O accesses are assumed to be memory-mapped) and to provide access protection on a block or page basis. Some protection is also available even if translation is disabled. Typically, it is not programmable. The AltiVec ISA does not provide any additional instructions to the PowerPC memory management model, but AltiVec instructions have options to ensure that an operand is correctly placed in a vector register or in memory.



## Chapter 2

# AltiVec Register Set

This chapter describes the register organization defined by AltiVec technology. It also describes how AltiVec instructions affect some of the registers in the PowerPC architecture. AltiVec Instruction Set Architecture (ISA) defines register-to-register operations for all computational instructions. Source data for these instructions is accessed from the on-chip vector registers (VRs) or are provided as immediate values embedded in the opcode. Architecturally, the VRs are separate from the general-purpose registers (GPRs) and floating-point registers (FPRs). Data is transferred between memory and vector registers with explicit AltiVec load and store instructions only.

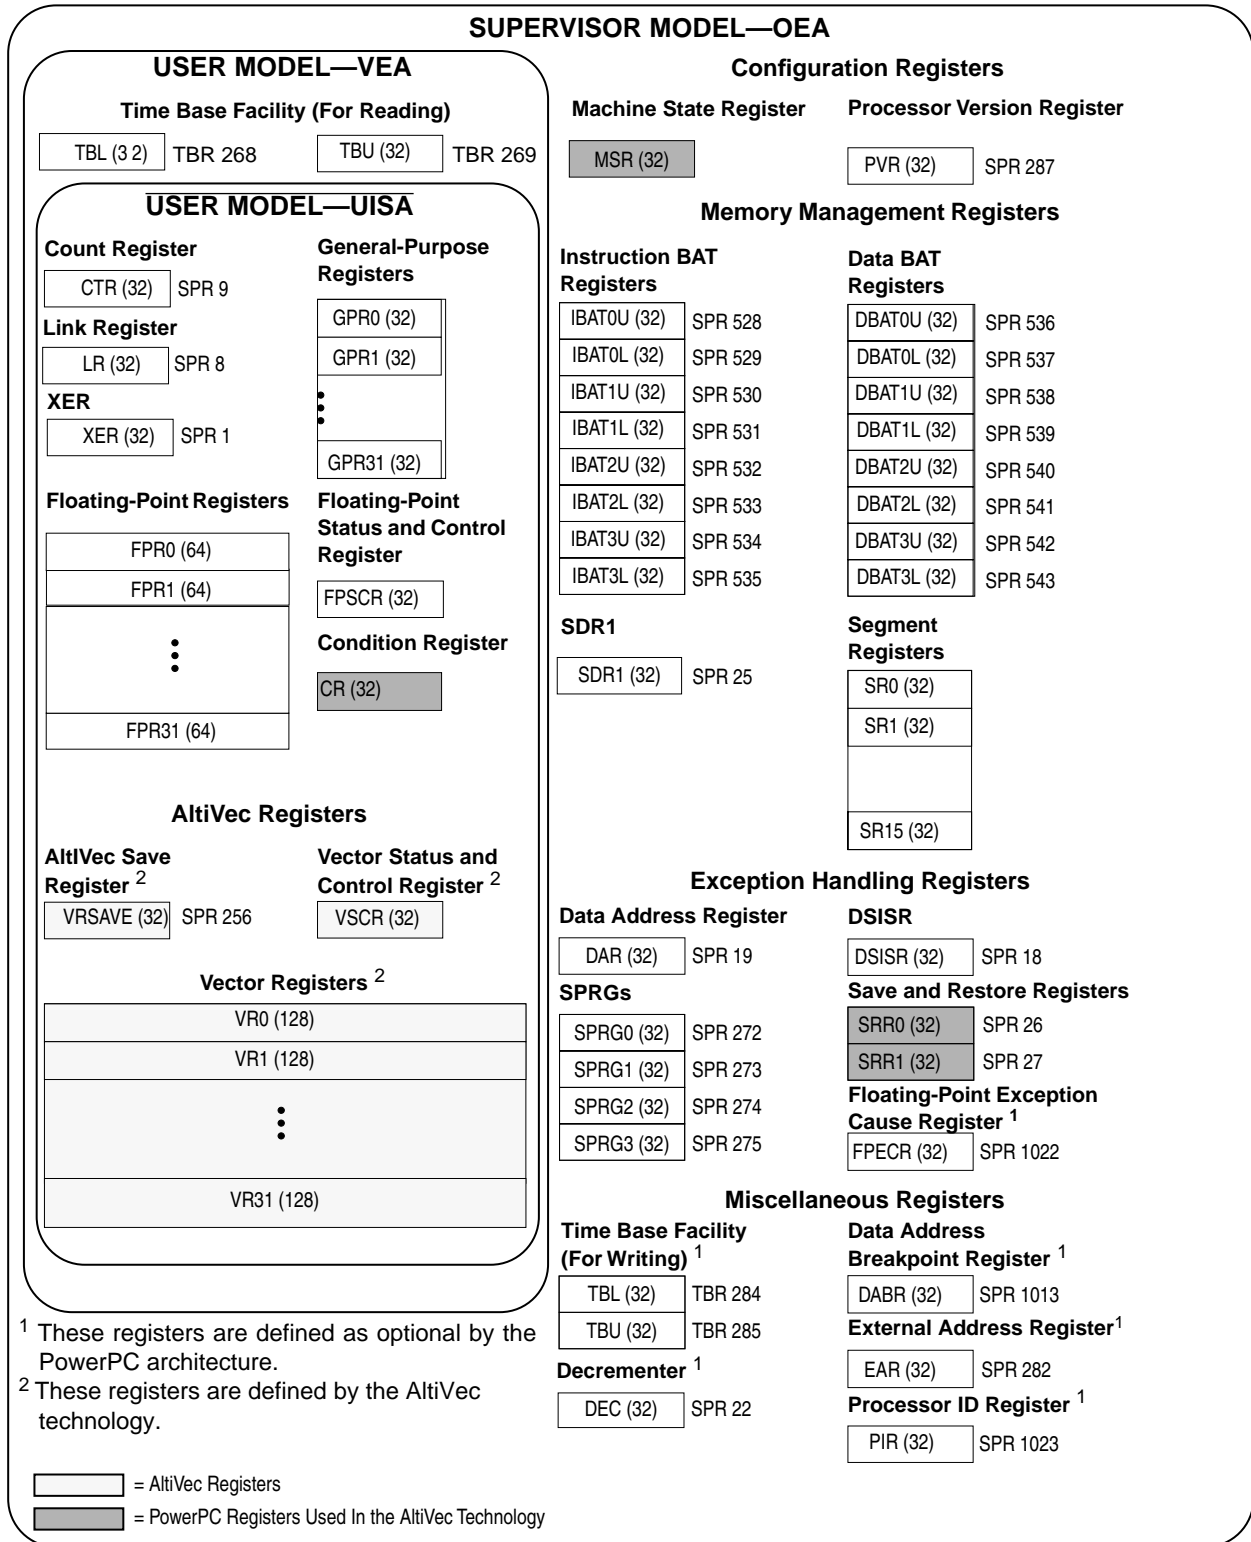
Note that the handling of reserved bits in any register is implementation-dependent. Software is permitted to write any value to a reserved bit in a register. However, a subsequent reading of the reserved bit returns 0 if the value last written to the bit was 0 and returns an undefined value (may be 0 or 1) otherwise. This means that even if the last value written to a reserved bit was 1, reading that bit may return 0.

### 2.1 Overview on the AltiVec and PowerPC Registers

The addition of AltiVec technology adds some additional new registers as well as affecting bit settings in some of the PowerPC registers when AltiVec instructions are executed. Figure 2-1 shows a graphic representation of the entire PowerPC register set and how the AltiVec register set resides within the PowerPC architecture. The PowerPC registers affected by AltiVec instructions are shaded and AltiVec registers are highlighted as well. Note that a processor that implements the PowerPC architecture may have additional registers specific only to that processor.

# Freescale Semiconductor, Inc.

## Overview on the AltiVec and PowerPC Registers



## 2.2 AltiVec Register Set Overview

AltiVec registers, shown in Figure 2-2 can be accessed by user or supervisor-level instructions. The vector registers (VRs) are accessed as instruction operands. Access to the registers can be explicit (that is, through the use of specific instructions for that purpose such as Move from Vector Status and Control Register (**mfvscr**) and Move to Vector Status and Control Register (**mtvscr**) instructions) or implicit as part of the execution of an instruction. The VRs are accessed both explicitly and implicitly.

The number to the right of the register name indicates the number used in the syntax of the instruction operands to access the register (for example, the number used to access the VRSAVE is SPR 256).

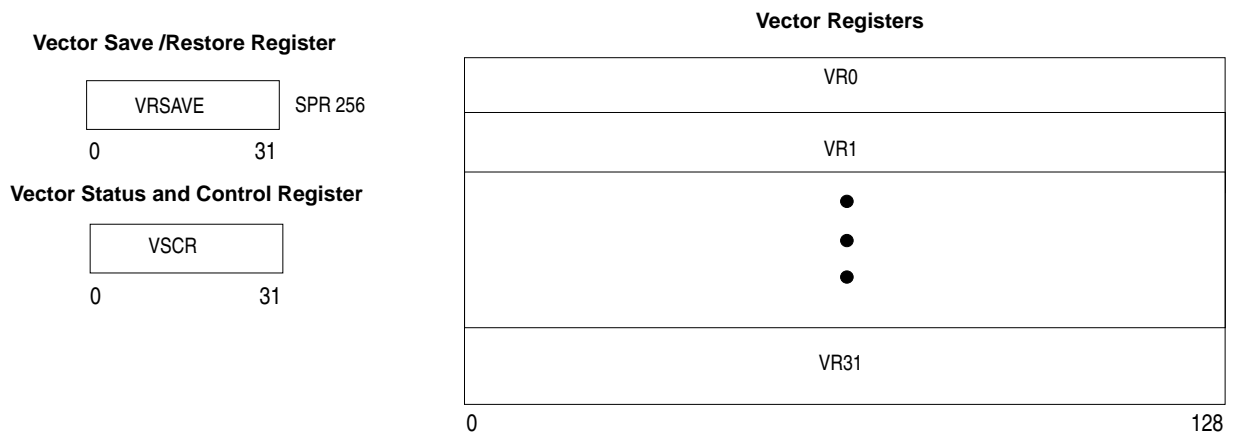


Figure 2-2. AltiVec Register Set

The user-level registers can be accessed by all software with either user or supervisor privileges. The user-level register set for AltiVec technology includes the following:

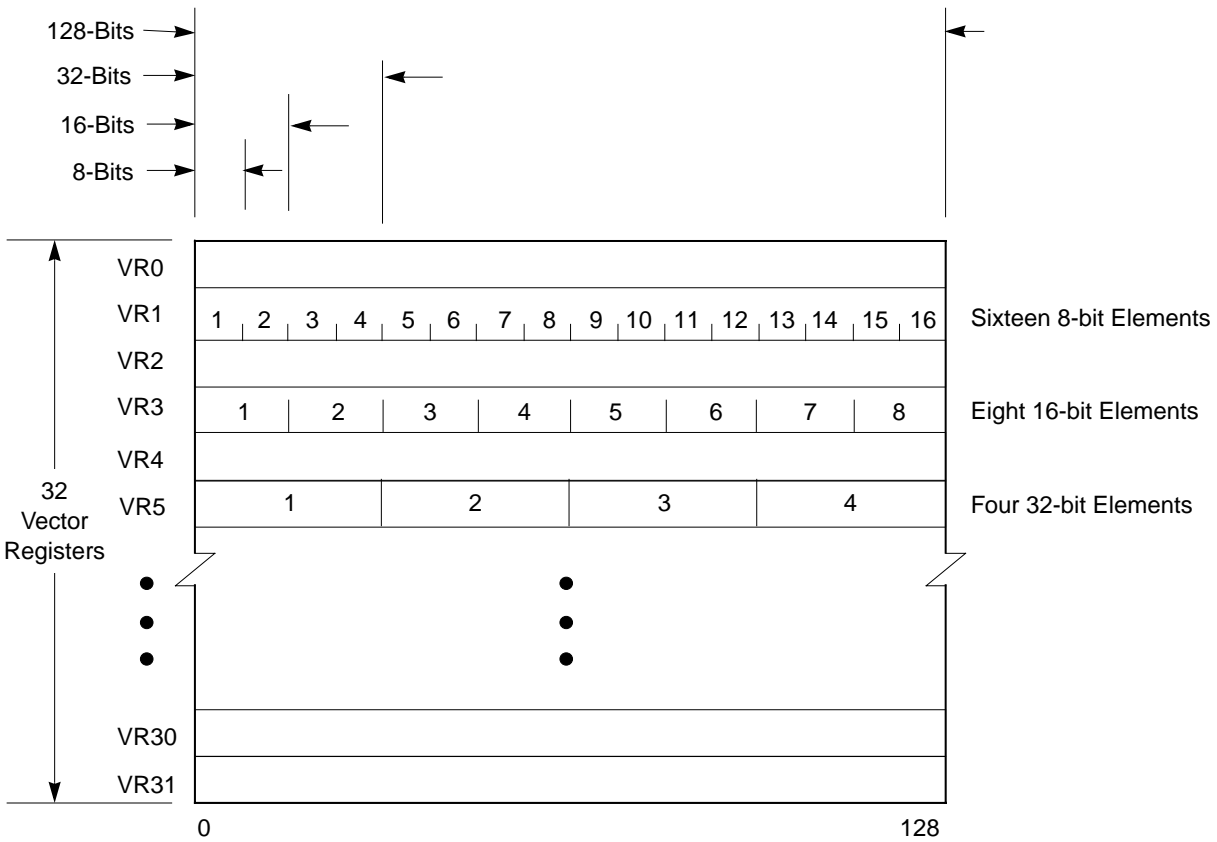
- **Vector registers (VRs):** The vector register file consists of 32 VRs designated as VR0–VR31. The VRs serve as vector source and vector destination registers for all vector instructions. See Section 2.3.2, “Vector Status and Control Register (VSCR),” for more information.
- **Vector status and control register (VSCR):** The VSCR contains the non-Java and saturation bit with the remaining bits being reserved. See Section 2.3.2, “Vector Status and Control Register (VSCR),” for more details.
- **Vector save/restore register (VRSAVE):** The VRSAVE assists the application and operating system software in saving and restoring the architectural state across context-switched events. The bits in the VRSAVE can indicate whether the vector register is live (1) or dead (0). See Section 2.3.3, “Vector Save/Restore Register (VRSAVE),” for more information.

## 2.3 Registers defined by AltiVec ISA

AltiVec ISA has defined several registers. The new AltiVec registers for the most part only interact with AltiVec instructions, with the exception of the VRSAVE register that is read or written by the PowerPC instructions **mfspr** or **mtspr**, respectively.

### 2.3.1 AltiVec Vector Register File (VRF)

The VRF, shown in Figure 2-3, has 32 registers, each 128 bits wide. Each vector register can hold sixteen 8-bit elements, eight 16-bit elements, or four 32-bit elements.




**Figure 2-3. Vector Registers (VRs)**

The vector registers are accessed as vector instruction operands. Access to registers are explicit as part of the execution of an AltiVec instruction.

### 2.3.2 Vector Status and Control Register (VSCR)

The vector status and control register (VSCR) is a 32-bit vector register (not an SPR) that is read and written in a manner similar to the FPSCR in the PowerPC scalar floating-point unit. The VSCR is shown in Figure 2-4



**Figure 2-4. Vector Status and Control Register (VSCR)**

The VSCR has two defined bits, the AltiVec non-Java mode (NJ) bit (VSCR[15]) and the AltiVec saturation (SAT) bit (VSCR[31]); the remaining bits are reserved.

Special instructions Move from Vector Status and Control Register (**mfvscr**) and Move to Vector Status and Control Register (**mtvscr**) are provided to move the contents of VSCR from and to a vector register. When moved to or from a vector register, the 32-bit VSCR is right-justified in the 128-bit vector register. When moved to a vector register, the upper 96 bits VR<sub>n</sub> [0–95] of the vector register are cleared, so the VSCR in a vector register looks as shown in Figure 2-5



**Figure 2-5. 32-bit VSCR Moved to a 128-bit Vector Register**

VSCR bit settings are shown in Table 2-1.

**Table 2-1. VSCR Field Descriptions**

Bit	Name	Description
0–14	—	Reserved. The handling of reserved bits is the same as that for other PowerPC registers. Software is permitted to write any value to such a bit. A subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value (0 or 1) otherwise.
15	NJ	Non-Java. This bit determines whether AltiVec floating-point operations are performed in a Java-IEEE-C9X-compliant mode or a possibly faster non-Java/non-IEEE mode. 0 The Java-IEEE-C9X-compliant mode is selected. Denormalized values are handled as specified by Java, IEEE, and the C9X standard. 1 The non-Java/non-IEEE-compliant mode is selected. If an element in a source vector register contains a denormalized value, the value 0 is used instead. If an instruction causes an underflow exception, the corresponding element in the target VR is cleared to 0. In both cases the 0 has the same sign as the denormalized or underflowing value. This mode is described in detail in the floating-point overview Section 3.2.1, “Floating-Point Modes.”

**Table 2-1. VSCR Field Descriptions (continued)**

Bit	Name	Description
16–30	—	Reserved. The handling of reserved bits is the same as that for other PowerPC registers. Software is permitted to write any value to such a bit. A subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value (0 or 1) otherwise.
31	SAT	<p>Saturation.</p> <p>A sticky status bit indicating that some field in a saturating instruction saturated since the last time SAT was cleared. In other words, when SAT = 1 it remains set to 1 until it is cleared to 0 by an <b>mtvscr</b> instruction. For further discussion refer to Section 4.2.1.1, “Saturation Detection.”</p> <p>0 Indicates no saturation occurred; <b>mtvscr</b> can explicitly clear this bit.</p> <p>1 The AltiVec saturate instruction is set when saturation occurs for the results one of AltiVec instructions having saturate in its name as follows:</p> <ul style="list-style-type: none"> <li>Move to VSCR (<b>mtvscr</b>)</li> <li>Vector Add Integer with Saturation (<b>vaddubs</b>, <b>vadduhs</b>, <b>vadduws</b>, <b>vaddsbs</b>, <b>vaddshs</b>, <b>vaddsws</b>)</li> <li>Vector Subtract Integer with Saturation (<b>vsububs</b>, <b>vsubuhs</b>, <b>vsubuws</b>, <b>vsubsbs</b>, <b>vsubshs</b>, <b>vsubsws</b>)</li> <li>Vector Multiply-Add Integer with Saturation (<b>vmhaddshs</b>, <b>vmhraddshs</b>)</li> <li>Vector Multiply-Sum with Saturation (<b>vmsumuhs</b>, <b>vmsumshs</b>, <b>vmsumsws</b>)</li> <li>Vector Sum-Across with Saturation (<b>vsumsws</b>, <b>vsum2sws</b>, <b>vsum4sbs</b>, <b>vsum4shs</b>, <b>vsum4ubs</b>)</li> <li>Vector Pack with Saturation (<b>vpkuhus</b>, <b>vpkuwus</b>, <b>vpkshus</b>, <b>vpkswus</b>, <b>vpkshss</b>, <b>vpkswss</b>)</li> <li>Vector Convert to Fixed-Point with Saturation (<b>vctuxs</b>, <b>vctxsx</b>)</li> </ul>

The **mtvscr** is context synchronizing. This implies that all AltiVec instructions logically preceding an **mtvscr** in the program flow execute in the architectural context (NJ mode) that existed before completion of **mtvscr**, and that all instructions logically following after **mtvscr** execute in the new context (NJ mode) established by the **mtvscr**.

After an **mfvscr** instruction executes, the result in the target vector register is architecturally precise. That is, it reflects all updates to the SAT bit that could have been made by vector instructions logically preceding it in the program flow, and further, it will not reflect any SAT updates that may be made to it by vector instructions logically following it in the program flow. Because it is context synchronizing, **mfvscr** can be much slower than typical AltiVec instructions, and therefore care must be taken in reading it to avoid performance problems.

### 2.3.3 Vector Save/Restore Register (VRSAVE)

The VRSAVE register shown in Figure 2-6 is a user-level 32-bit SPR used to assist in application and operating system software in saving and restoring the architectural state across process context-switched events. The VRSAVE is SPR 256 and is entirely maintained and managed by software.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Field	VR0	VR1	VR2	VR3	VR4	VR5	VR6	VR7	VR8	VR9	VR10	VR11	VR12	VR13	VR14	VR15
Reset	0000_0000_0000_0000															
R/W	R/W with <b>mf spr</b> or <b>mt spr</b> instruction															

	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Field	VR16	VR17	VR18	VR19	VR20	VR21	VR22	VR23	VR24	VR25	VR26	VR27	VR28	VR29	VR30	VR31
Reset	0000_0000_0000_0000															
R/W	R/W with <b>mf spr</b> or <b>mt spr</b> instructions															
SPR	SPR256															

Figure 2-6. Vector Save/Restore Register (VRSAVE)

VRSAVE bit settings are shown in Figure 2-2

Table 2-2. VRSAVE Bit Settings

Bits	Name	Description
0-31	VR $n$	Each bit in the VRSAVE register indicates whether the corresponding VR contains data in use by the executing process. 0 VR $n$ is not being used for the current process 1 VR $n$ is using VR $n$ for the current process

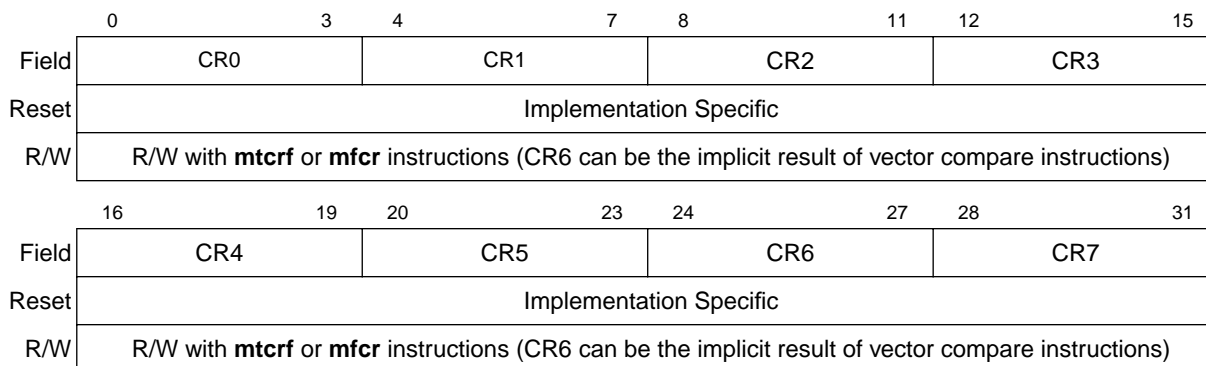
The VRSAVE register can be accessed only by the **mf spr** and **mt spr** instructions. Each bit in this register corresponds to a vector register (VR) and indicates whether the corresponding register contains data that is currently in use by the executing process. Therefore, the operating system needs to save and restore only those VRs when an exception occurs. If this approach is taken, it must be applied rigorously; if a program fails to indicate that a given VR is in use, software errors may occur that are difficult to detect and correct because they are timing-dependent. Some operating systems save and restore VRSAVE only for programs that also use other AltiVec registers.

## 2.4 Additions to PowerPC UISA Registers

The PowerPC UISA registers can be accessed by either user- or supervisor-level instruction. The one register affected by AltiVec architecture is the condition register (CR). The CR is a 32-bit register, divided into eight 4-bit fields, CR0–CR7, that reflects the results of certain arithmetic operations and provides a mechanism for testing and branching. For more details refer to Chapter 2, “Register Set,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

## 2.4.1 PowerPC Condition Register

The PowerPC condition register (CR) is a 32-bit register that reflects the result of certain operations and provides a mechanism for testing and branching. For AltiVec ISA, the CR6 field can optionally be used, that is if an AltiVec instruction field's record bit (Rc) is set in a vector compare instruction. The CR6 field is updated. The CR is divided into eight 4-bit fields, CR0–CR7, as shown in Figure 2-7



**Figure 2-7. Condition Register (CR)**

For more details on the CR see Chapter 2, “Register Set,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

To control program flow based on vector data, all vector compare instructions can optionally update CR6. If the instruction field's record bit (Rc) is set in a vector compare instruction, CR6 is updated according to Table 2-3.

**Table 2-3. CR6 Field's Bit Settings for Vector Compare Instructions**

CR Bit	CR6 Field Bit	Vector Compare	Vector Compare Bounds
24	0	1 Relation is true for all element pairs	0
25	1	0	0
26	2	1 Relation is false for all element pairs 0 All fields were in bounds	1 All fields are in bounds for the <b>vcmpbfp</b> instruction so the result code of all fields is 0b00 0 One of the fields is out of bounds for the <b>vcmpbfp</b> instruction
27	3	0	0

The Rc bit should be used sparingly because when Rc = 1 it can cause a somewhat longer latency or be more disruptive to instruction pipeline flow than when Rc = 0. Therefore techniques of accumulating results and testing infrequently are advised.

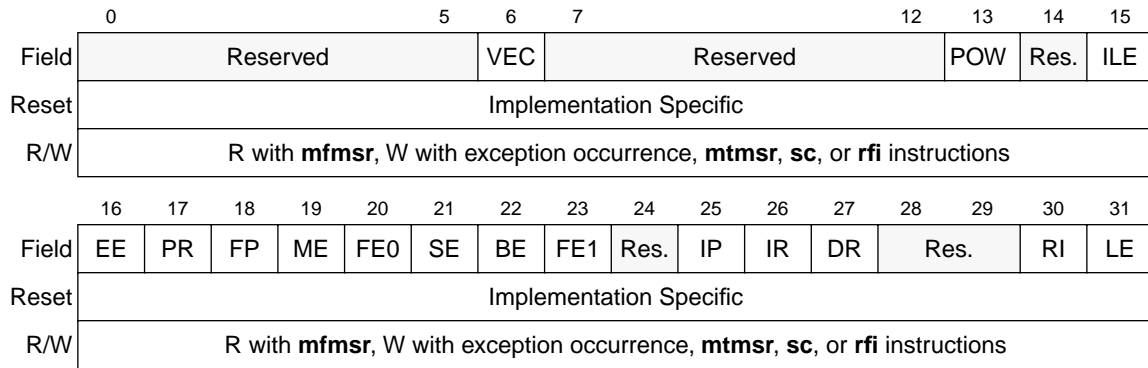


## 2.5 Additions to PowerPC OEA Registers

The PowerPC operating environment architecture (OEA) can be accessed only by supervisor-level instructions. Any attempt to access these SPRs with user-level instructions results in a supervisor-level exception. For more details on the MSR and SRR see Chapter 2, “Register Set,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 2.5.1 AltiVec Field added in the PowerPC Machine State Register (MSR)

An AltiVec available field is added to the PowerPC machine state register (MSR). The MSR is 32 bits wide as shown in Figure 2-8.



**Figure 2-8. Machine State Register (MSR)**

In 32-bit PowerPC implementations, bit 6, the VEC field, is added to the MSR as shown in Figure 2-8. Also AltiVec data stream prefetching instructions will be suspended and resumed based on MSR[PR] and MSR[DR]. The Data Stream Touch (**dst**) and Data Stream Touch for Store (**dstst**) instructions are supported whenever MSR[DR] = 1. If either instruction is executed when MSR[DR] = 0 (real addressing mode), the results are boundedly undefined. For each existing data stream, prefetching is enabled if MSR[DR] = 1 and MSR[PR] has the value it had when the **dst** or **dstst** instruction that specified the data stream was executed. Otherwise prefetching for the data stream is suspended. In particular, the occurrence of an exception suspends all data stream prefetching.

Table 2-4 shows AltiVec bit definitions for the MSR as well as how the PR and DR bits are affected by AltiVec data stream instructions.

**Table 2-4. MSR Bit Settings**

Bits	Name	Description
6	VEC	AltiVec Available 0 AltiVec is disabled. 1 AltiVec is enabled. Note: Any attempt to execute a non-stream AltiVec instruction when the bit is cleared causes the processor to execute an “AltiVec Unavailable Exception” when the instruction accesses the VRF or VSCR register. This exception does not happen for data streaming instructions ( <b>dst(t)</b> , <b>dstst(t)</b> , and <b>dss</b> ), that is, the VRF and VSCR registers are available to the data streaming instructions even when the MSR[VEC] is cleared. The VRSAVE register is not protected by MSR [VEC], that is, it can be accessed even when MSR[VEC] is cleared.
17	PR	Privilege level 0 The processor can execute both user- and supervisor-level instructions. 1 The processor can only execute user-level instructions.  Note: Care should be taken if data stream prefetching is used in supervisor mode (MSR[PR] = 0). For each existing data stream, prefetching is enabled if MSR[DR] = 1 and MSR[PR] has the value it had when the <b>dst</b> or <b>dstst</b> instruction that specified the data stream was executed. Otherwise prefetching for the data stream is suspended.
27	DR	Data address translation 0 Data address translation is disabled. If data stream touch ( <b>dst</b> ) and data stream touch for store ( <b>dstst</b> ) instructions are executed whenever DR = 0, the results are boundedly undefined 1 Data address translation is enabled. Data stream touch ( <b>dst</b> ) and data stream touch for store ( <b>dstst</b> ) instructions are supported whenever DR = 1.


For more detailed information including the other bit settings for MSR, refer to Chapter 2, “Register Set,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

## 2.5.2 Machine Status Save/Restore Registers (SRRs)

The machine status save/restore registers (SRRs) are part of the PowerPC OEA supervisor-level registers. The SRR0 and SRR1 registers are used to save machine status on exceptions and to restore machine status when an **rfi** instruction is executed. For more detailed information, refer to Chapter 2, “Register Set,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 2.5.2.1 Machine Status Save/Restore Register 0 (SRR0)


The SRR0 is a 32-bit register in 32-bit implementation. SRR0 is used to save machine status on exceptions and restore machine status when an **rfi** instruction is executed. For AltiVec ISA, it holds the effective address (EA) for the instruction that caused the AltiVec unavailable exception. The AltiVec unavailable exception occurs when no higher priority exception exists, and an attempt is made to execute an AltiVec instruction when MSR[VEC] = 0. The format of SRR0 is shown in Figure 2-9.



**Figure 2-9. Machine Status Save/Restore Register 0 (SRR0)**

### 2.5.2.2 Machine Status Save/Restore Register 1 (SRR1)

The SRR1 is a 32-bit register in 32-bit implementation. SRR1 is used to save machine status on exceptions and to restore machine status when an *rfi* instruction is executed. The format of SRR1 is shown in Figure 2-10.



**Figure 2-10. Machine Status Save/Restore Register 0 (SRR1)**

When an AltiVec unavailable exception occurs, SRR1[1–4] and SRR[10–15] are cleared and all other SRR1 bits are loaded from the MSR as it was just prior to the interrupt. So MSR[0], MSR[5–9], and MSR[16–31] are placed into the corresponding bit positions of SRR1 as they were before the exception was taken.



# Chapter 3

## Operand Conventions

This chapter describes the operand conventions as they are represented in AltiVec technology at the user instruction set architecture (UISA) level. Detailed descriptions are provided of conventions used for transferring data between vector registers and memory, and representing data in these vector registers using both big- and little-endian byte ordering. Additionally, the floating-point default conditions for exceptions are described.

### 3.1 Data Organization in Memory U

In addition to supporting byte, half-word and word operands, as defined in the PowerPC architecture UISA, AltiVec instruction set architecture (ISA) supports quad-word (128-bit) operands.

The following sections describe the concepts of alignment and byte ordering of data for quad words, otherwise alignment is the same as described in Chapter 3, “Operand Conventions,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

#### 3.1.1 Aligned and Misaligned Accesses

Vectors are accessed from memory with instructions such as Vector Load Indexed (**lvx**) and Store Vector Indexed (**stvx**) instructions. The operand of a vector register to memory access instruction has a natural alignment boundary equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. Each AltiVec instruction is a 4-byte word and is word-aligned like PowerPC instructions.

Operands for vector register to memory access instructions have the characteristics shown in Table 3-1.

**Table 3-1. Memory Operand Alignment**

Operand	Length	32-bit Aligned Address (28–31) <sup>1</sup>
Byte	8 bits (1 byte)	xxxx
Half word	2 bytes	xxx0
Word	4 bytes	xx00
Quad word	16 bytes	0000

<sup>1</sup> An x in an address bit position indicates that the bit can be 0 or 1 independent of the state of other bits in the address

The concept of alignment is also applied more generally to data in memory. For example, an 8-byte data item is said to be half-word-aligned if its address is a multiple of two; that is, the effective address (EA) points to the next effective address that is 2 bytes (a half word) past the current effective address (EA + 2 bytes), and then the next being the EA + 4 bytes, and effective address would continue skipping every 2 bytes (2 bytes = 1 half word). This ensures that the effective address is half-word aligned as it points to each successive half word in memory.

It is important to understand that AltiVec memory operands are assumed to be aligned, and AltiVec memory accesses are performed as if the appropriate number of low-order bits of the specified effective address were zero. This assumption is different from PowerPC integer and floating-point memory access instructions where alignment is not always assumed. So for AltiVec ISA, the low-order bit of the effective address is ignored for half-word AltiVec memory access instructions, and the low-order four bits of the effective address are ignored for quad-word AltiVec memory access instructions. The effect is to load or store the memory operand of the specified length that contains the byte addressed by the effective address.

If a memory operand is misaligned, additional instructions must be used to correctly place the operand in a vector register or in memory. AltiVec technology provides instructions to shift and merge the contents of two vector registers. These instructions facilitate copying misaligned quad-word operands between memory and the vector registers.

### 3.1.2 AltiVec Byte Ordering

For processors that implement the PowerPC architecture and AltiVec technology, the smallest addressable memory unit is the byte (8 bits), and scalars are composed of one or more sequential bytes. AltiVec ISA supports both big- and little-endian byte ordering. The default byte ordering is big-endian. However, the code sequence used to switch from big- to little-endian mode may differ among processors.

The PowerPC architecture uses the machine state register (MSR) for specifying byte ordering in little-endian mode (LE). A value of 0 specifies big-endian mode and a value of 1 specifies little-endian mode. For further details on byte ordering in the PowerPC architecture, refer to Chapter 3, “Operand Conventions,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

AltiVec ISA follows the endian support of the PowerPC architecture for elements up to double words with additional support for quad words. In AltiVec ISA when a 64-bit scalar is moved from a register to memory, it occupies eight consecutive bytes in memory and a decision must be made regarding byte ordering in these eight addresses.

### 3.1.2.1 Big-Endian Byte Ordering

For big-endian scalars, the most-significant byte (MSB) is stored at the lowest (or starting) address while the least-significant byte (LSB) is stored at the highest (or ending) address. This is called big-endian because the big end of the scalar comes first in memory.

### 3.1.2.2 Little-Endian Byte Ordering

For little-endian scalars, the LSB is stored at the lowest (or starting) address while the MSB is stored at the highest (or ending) address. This is called little-endian because the little end of the scalar comes first in memory.

### 3.1.3 Quad Word Byte Ordering Example

The idea of big- and little-endian byte ordering is best illustrated in an example of a quad word such as 0x0011\_2233\_4455\_6677\_8899\_AA\_BB\_CCDD\_EEFF located in memory. This quad word is used throughout this section to demonstrate how the bytes that comprise a quad word are mapped into memory.

The quad word (0x0011\_2233\_4455\_6677\_8899\_AA\_BB\_CCDD\_EEFF) is shown in big-endian mapping in Figure 3-1. A hexadecimal representation is used for showing address values and the values in the contents of each byte. The address is shown below each byte’s contents. The big-endian model addresses the quad word at address 0x00, which is the MSB (0x00), proceeding to the address 0x0F, which contains the LSB (0xFF)

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word															
Contents	00	11	22	33	44	55	66	77	88	99	AA	BB	CC	DD	EE	FF
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	↑															↑
	MSB															LSB

**Figure 3-1. Big-Endian Mapping of a Quad Word**

Figure 3-2 shows the same quad word using little-endian mapping. In the little-endian model, the quad word's 0x00 address specifies the LSB (0xFF) and proceeds to address 0x0F which contains its MSB (0x00).

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15																
	Quad Word																															
Contents	FF	EE	DD	CC	BB	AA	99	88	77	66	55	44	33	22	11	00																
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F																
	↑															↑																
	LSB																MSB															

**Figure 3-2. Little-Endian Mapping of a Quad Word**

Figure 3-2 shows the sequence of bytes laid out with addresses increasing from left to right. Programmers familiar with little-endian byte ordering may be more accustomed to viewing quad words laid out with addresses increasing from right to left, as shown in Figure 3-3.

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15																
	Quad Word																															
Contents	00	11	22	33	44	55	66	77	88	99	AA	BB	CC	DD	EE	FF																
Address	0F	0E	0D	0C	0B	0A	09	08	07	06	05	04	03	02	01	00																
	↑															↑																
	MSB																LSB															

**Figure 3-3. Little-Endian Mapping of Quad Word—Alternate View**

This allows the little-endian programmer to view each scalar in its natural byte order of MSB to LSB. This section uses both conventions based on ease of understanding for the specific example.

### 3.1.4 Aligned Scalars in Little-Endian Mode

The effective address (EA) calculation for the load and store instructions is described in Chapter 4, “Addressing Modes and Instruction Set Summary.” For processors that implement the PowerPC architecture in little-endian mode, the effective address is modified before being used to access memory. In the PowerPC architecture, the three low-order address bits of the effective address are exclusive-ORed (XOR) with a three-bit value that depends on the length of the operand (1, 2, 4, or 8 bytes), as shown in Table 3-2. This address modification is called munging.



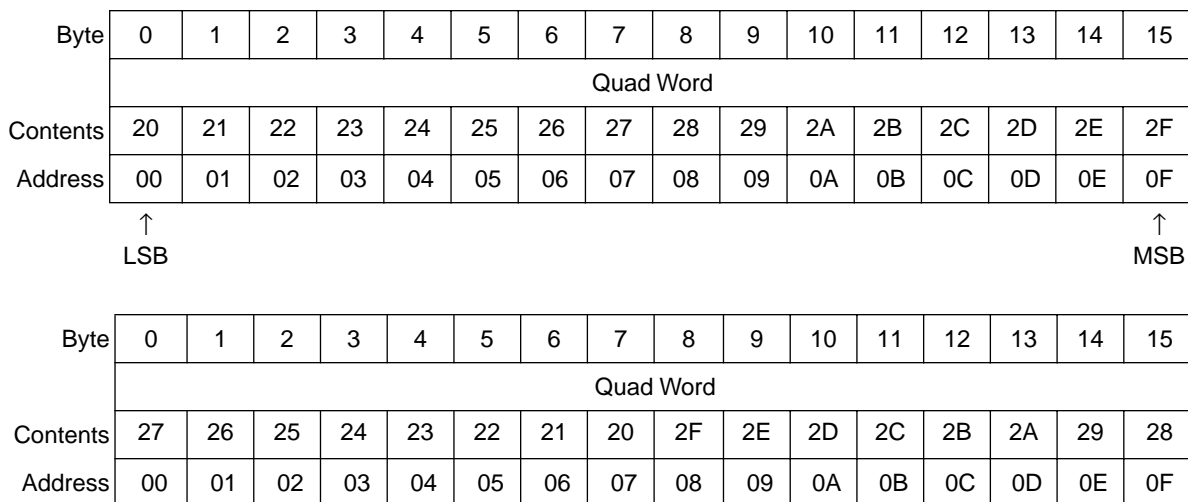
**Table 3-2. Effective Address Modifications**

Data Width (Bytes)	EA Modification
1	XOR with 0b111
2	XOR with 0b110
4	XOR with 0b100
8	No change

The munged physical address is passed to the cache or to main memory, and the specified width of the data is transferred (in big-endian order—that is, MSB at the lowest address, LSB at the highest address) between a GPR or FPR and the addressed memory locations (as modified).

Munging makes it appear to the processor that individual aligned scalars are stored as little-endian, when in fact they are stored in big-endian order but at different byte addresses within double words. Only the address is modified, not the byte order. For further details on how to align scalars in little-endian mode see Chapter 3, “Operand Conventions,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

The PowerPC address munging is performed on double-word units. In the PowerPC architecture, little-endian mode would have the double words of a quad word appear swapped. When the quad word in memory shown at the top of Figure 3-4, loads from address 0x00, the bottom of Figure 3-4 shows how it appears to the processor as it munges the address.



**Figure 3-4. Quad Word Load with PowerPC Munged Little-Endian Applied**

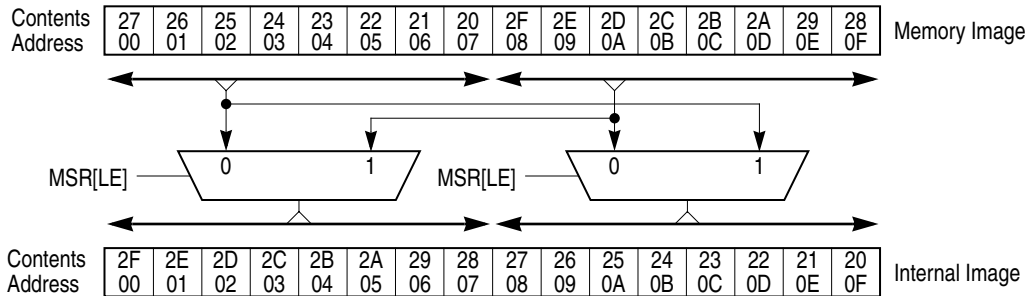
Note that double words are swapped. The byte element addressed by the quad word’s base address, 0x0F, contains 0x28, while its MSB at address 0x00 contains 0x27. This is due to the PowerPC munging being applied to offsets within double words; AltiVec ISA requires a munge within quad words.

## Freescale Semiconductor, Inc.

### Data Organization in Memory

To accommodate the quad-word operands, the PowerPC architecture cannot simply be extended by munging an extra address bit. It would break existing code or platforms. Processors that implement AltiVec technology could not be mixed with non-AltiVec processors. Instead, AltiVec processors implement a double-word swap when moving quad words between vector registers and memory.

Figure 3-5 shows how this swapping could be implemented. This diagram represents the load path double-word swapping; the store path looks the same, except that the memory and internal boxes are reversed.



**Figure 3-5. AltiVec Little Endian Double-Word Swap**

In the diagram, the numbers at the bottom of the byte boxes represent the offset address of that byte; the numbers at the top are the values of the bytes at that offset. The little-endian ordering is discontinuous because the PowerPC munging is performed only on double-word units. The purpose of the double word swap within the AltiVec unit is to perform an additional swap that is not part of the PowerPC architecture.

When MSR[LE] = 1, double words are swapped and the bytes appear in their expected ordering. When MSR[LE] = 0, no swapping occurs.

To summarize, in little-endian mode, the load vector element indexed instructions (**lvebx**, **lvehx**, and **lvewx**) and the store vector element indexed instructions (**stvebx**, **stvehx**, and **stvewx**) have the same 3-bit address munge applied to the memory address as is specified by the PowerPC architecture for integer and floating-point loads and stores. For the quad word load vector indexed instructions (**lvx** and **lvxl**) and the store vector indexed instructions (**stvx**, **stvxl**), the two double words of the quad-word scalar data are munged and swapped as they are moved between the vector register and memory.

### 3.1.5 Vector Register and Memory Access Alignment

When loading an aligned byte, half word, or word memory operand into a vector register, the element that receives the data is the element that would have received the data had the entire aligned quad word containing the memory operand addressed by the effective address been loaded. Similarly, when an element in a vector register is stored into an aligned memory operand, the element selected to be stored is the element that would have been stored into the memory operand addressed by the effective address had the entire

vector register been stored to the aligned quad word containing the memory operand addressed by the effective address. The position of the element in the target or source vector register depends on the endian mode, as described above. (Byte memory operands are always aligned.)

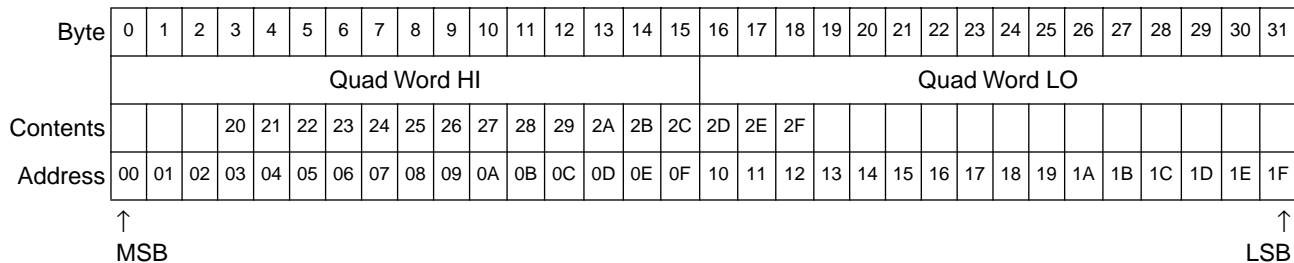
For aligned byte, half word, and word memory operands, if the corresponding element number is known when the program is written, the appropriate vector splat and vector permute instructions can be used to copy or replicate the data contained in the memory operand after loading the operand into a vector register. Vector splat instructions will take the contents of an element in a vector register and replicates them into each element in the destination vector register. A vector permute instruction is the concatenation of the contents of two vectors. An example of this is given in detail in Section 3.1.6, “Quad-Word Data Alignment.” Another method is to replicate the element across an entire vector register before storing it into an arbitrary aligned memory operand of the same length; the replication ensures that the correct data is stored regardless of the offset of the memory operand in its aligned quad word in memory.

Because vector loads and stores are size-aligned, application binary interfaces (ABIs) should specify, and programmers should take care to align data on quad-word boundaries for maximum performance.

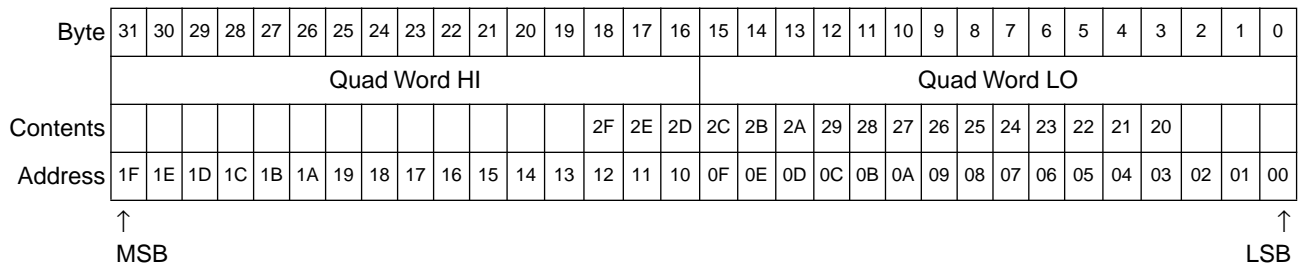
### 3.1.6 Quad-Word Data Alignment

AltiVec ISA does not provide for alignment exceptions for loading and storing data. When performing vector loads and stores, the effect is as if the low-order four bits of the address are 0x0, regardless of the actual effective address generated. Because vectors may often be misaligned due to the nature of the algorithm, AltiVec ISA provides support for post-alignment of quad-word loads and pre-alignment for quad-word stores. Note that in the following diagrams, the effect of the swapping described above is assumed and the memory diagrams will be shown with respect to the logical mapping of the data.

Figure 3-6 and Figure 3-7 show misaligned vectors in memory for both big- and little-endian ordering. The big-endian and little-endian examples assumes that the desired vector begins at address 0x03. In the figure, HI denotes high-order quad word, and LO means low-order quad word.



**Figure 3-6. Misaligned Vector in Big-Endian Mode**




**Figure 3-7. Misaligned Vector in Little-Endian Addressing Mode**

Figure 3-6 and Figure 3-7 show how such misaligned data causes data to be split across aligned quad words; only aligned quad words are loaded or stored by AltiVec load/store instructions. To align this vector, a program must load both (aligned) quad words that contain a portion of the misaligned vector data and then execute a Vector Permute (`vperm`) instruction to align the result.

### 3.1.6.1 Accessing a Misaligned Quad Word in Big-Endian Mode

Figure 3-1 shows the big-endian alignment model. Using the example in Figure 3-8, `vHI` and `vLO` represent vector registers that contain the misaligned quad words containing the MSBs and LSBs, respectively, of the misaligned quad word; `vD` is the target vector register.



**Figure 3-8. Big-Endian Quad Word Alignment**

Alignment is performed by left-rotating the combined 32-byte quantity (`vHI:vLO`) by an amount determined by the address of the first byte of the desired data. This left-rotation is done by means of a `vperm` instruction whose control vector is generated by a Load Vector for Shift Left (`lvsl`) instruction after loading the most-significant quad word (MSQ) and least-significant quad word (LSQ) that contain the desired vector. The `lvsl` instruction uses the same address specification as the load vector indexed that loads the `vHI` component, which for big-endian ordering is the address of the desired vector.

The following instruction sequence extracts the quad word in big-endian mode:

```

lvx      vHI, rA, rB      # load the MSQ
lvsl     vP, rA, rB      # set the permute vector
addi     rB, rB, 16      # address of LSQ
lvx      vLO, rA, rB     # load LSQ component
vperm    vD, vHI, vLO, vP # align the data
    
```

Note that when data streaming is used, the overhead of generating the alignment permute vector can be spread out and the latency of the loads may be absorbed by using loop unrolling.

The process of storing a misaligned vector is essentially the reverse of that for loading, except that the code has a read-modify-write sequence. The logical algorithm is that the vector source must be right-shifted and split into two parts, each of which is merged (via a Vector Select (**vsel**) instruction) with the current contents of its MSQ and its LSQ and stored back using a Store Vector Indexed (**svx**) instruction.

The Load Vector for Shift Right (**lvsr**) instruction is used to produce the permute control vector to be used for the right-shifting. Note that a single register can be used for the shifted contents if a right-rotate is done. The rotate is performed by specifying the source register for both components of the Vector Permute (**vperm**); that is, a shift of a double register with the same contents in both parts results in a rotate. In addition, the same permute control vector can be used on a sequence of ones and zeros to generate a mask for use by the **vsel** instruction to do the merging.

The complete code sequence for the store case is as follows:

```

lvx      vHI, rA, rB          # load current MSQ for update
lvsr     vP, rA, rB          # load the alignment vector
addi     rB, rB, 16          # address of LSQ
lvx      vLO, rA, rB         # load the current LSQ's data
vspltisbv1s, -1             # generate the select mask bits
vspltisbv0s, 0
vperm    vMask, v0s, v1s, vP # right shift the select mask
vperm    vSrc, vSrc, vSrc, vP # right rotate the data
vsel     vLO, vSrc, vLO, vMask # insert LSQ component
vsel     vHI, vHI, vSrc, vMask # insert MSQ component
stvx     vLO, rA, rB         # store LSQ
addi     rB, rB, -16         # address of MSQ
stvx     vHI, rA, rB         # store MSQ

```

When fetching a misaligned stream, the control vector need only be computed once. Thus the time required for aligned fetches on the ends of the stream is proportioned out. None of the data fetched internally to the stream is wasted and only gets fetched once. The average time spent for a misaligned **lvx** instruction in a long sequence approaches the latency of one **lvx** and one **vperm** instruction.

### 3.1.6.2 Accessing a Misaligned Quad Word in Little-Endian Mode

The instruction sequences used to access misaligned quad-word operands in little-endian mode are similar to those used in big-endian mode. The following instruction sequence can be used to load the misaligned quad word shown in Figure 3-7 into a vector register in little-endian mode. The load alignment case is shown in Figure 3-9. The vector register **vHI** and **vLO** receive the **MSQ** and **LSQ** respectively; **vD** is the target vector register. The **lvsr** instruction uses the same address specification as an **lvx** that loads **vLO**; in little-endian byte ordering this is the address of the desired misaligned quad word.

```

lvx      vLO,rA,rB      # load the LSQ
lvsr     vP,rA,rB      # set the permute vector
addi     rB,rB,16      # address of MSQ
lvx      vHI,rA,rB     # load MSQ component
vperm    vD,vHI,vLO,vP # align the data
    
```

Similarly, the following sequence of instructions stores the contents of register **vD** into a misaligned quad word in memory in little-endian mode.

```

lvx      vLO,rA,rB     # load current LSQ for update
lvsl     vP,rA,rB     # load the alignment vector
addi     rB,rB,16     # address of MSQ
lvx      vHI,rA,rB    # load the current MSQ's data
vspltib  vls,-1      # generate the select mask bits
vspltib  v0s,0       #
vperm    vMask,v0s,vls,vP # left rotate the select mask
vperm    vSrc,vSrc,vSrc,vP # left rotate the data
vsel     vHI,vHI,vSrc,vMask # insert MSQ component
vsel     vLO,vSrc,vLO,vMask # insert LSQ component
stvx    vHI,rA,rB    # store MSQ
addi     rB,rB,-16    # address of LSQ
stvx    vLO,rA,rB    # store LSQ
    
```




Figure 3-9. Little-Endian Alignment

### 3.1.6.3 Scalar Loads and Stores

No alignment is performed for scalar load or store instructions in AltiVec ISA. If a vector load or store address is not properly size aligned, the suitable number of least significant bits are ignored and a size aligned transfer occurs instead. Data alignment must be performed explicitly after being brought into the registers. No assistance is provided for aligning individual scalar elements that are not aligned on their natural boundary. The placement of scalar data in a vector element depends upon its address. That is, the placement of the addressed scalar is the same as if a load vector indexed instruction has been performed, except that only the addressed scalar is accessed (for cache-inhibited space); the values in the other vector elements are boundedly undefined. Also, data in the specified scalar is the same as if a store vector indexed instruction had been performed, except that only the scalar addressed is affected. No instructions are provided to assist in aligning individual scalar elements that are not aligned on their natural size boundary.

When a program knows the location of a scalar, it can perform the correct vector splats and vector permutes to move data to where it is required. For example, if a scalar is to be used as a source for a vector multiply (that is, each element multiplied by the same value), the scalar must be splatted into a vector register. Likewise, a scalar stored to an arbitrary memory location must be splatted into a vector register, and that register must be specified as the source of the store. This guarantees that the data appears in all possible positions of that scalar size for the store.

### 3.1.6.4 Misaligned Scalar Loads and Stores

Although no direct support of misaligned scalars is provided, the load-aligning sequence for big-endian vectors described in Section 3.1.6.1, “Accessing a Misaligned Quad Word in Big-Endian Mode,” can be used to position the scalar to the left vector element, which can then be used as the source for a splat. That is, the address of a scalar is also the address of the left-most element of the quad word at that address. Similarly, the read-modify-write sequences, with the mask adjusted for the scalar size, can be used to store misaligned scalars. The same is true for little-endian mode, the load-aligning sequence for little-endian vectors described Section 3.1.6.2, “Accessing a Misaligned Quad Word in Little-Endian Mode” can be used to position the scalar to the right vector element, which can then be used

as the source for a splat. That is, the address of a scalar is also the address of the right-most element of the quad word at that address.

Note that while these sequences work in cache-inhibited space, the physical accesses are not guaranteed to be atomic.

### 3.1.7 Mixed-Endian Systems

In many systems, the memory model is not as simple as the examples in this chapter. In particular, big-endian systems with subordinate little-endian buses (such as PCI) comprise a mixed-endian environment.

The basic mechanism to handle this is to use the Vector Permute (**vperm**) instruction to swap bytes within data elements. The value of the permute control vector depends on the size of the elements (8, 16, 32). That is, the permute control vector performs a parallel equivalent of the Load Word Byte-Reverse Indexed (**lwbrx**) PowerPC instruction within the vector registers.

The ultimate problem occurs when there are misaligned, mixed-endian vectors. This can be handled by applying a vector permute of the data as required for the misaligned case, followed by the swapping vector permute on that result. Note that for streaming cases, the effect of this double permute can be accomplished by computing the swapping permute of the alignment permute vector and then applying the resulting permute control vector to incoming data.

## 3.2 AltiVec Floating-Point Instructions—UISA

■ There are two kinds of floating-point instructions defined for the PowerPC ISA and AltiVec ISA:

- computational
- noncomputational

Computational instructions are defined by the IEEE-754 standard for 32-bit arithmetic (those that perform addition, subtraction, multiplication, and division) and the multiply-add defined by the architecture. Noncomputational floating-point instructions consist of the floating-point load and store instructions. Only the computational instructions are considered floating-point operations throughout this chapter.

The single-precision format, value representations, and computational model to be defined in Chapter 3, “Operand Conventions,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*, apply to AltiVec floating-point except as follows:



- In general, no status bits are set to reflect the results of floating-point operations. The only exception is that VSCR[SAT] may be set by the Vector Convert to Fixed-Point Word instructions.
- With the exception of the two Vector Convert to Fixed-Point Word (**vctuxs**, **vctxsx**) instructions and three of the four Vector Round to Floating-Point Integer (**vrfiz**, **vrfig**, **vrfim**) instructions, all AltiVec floating-point instructions that round use the round-to-nearest rounding mode.
- Floating-point exceptions cannot cause the system error handler to be invoked.

If a function is required that is specified by the IEEE standard, is not supported by AltiVec ISA, and cannot be emulated satisfactorily using the functions that are supported by AltiVec ISA, the functions provided by the floating-point processor should be used; see Chapter 4, “Addressing Modes and Instruction Set Summary,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 3.2.1 Floating-Point Modes

AltiVec ISA supports two floating-point modes of operation—a Java mode and a non-Java mode of operation that is useful in circumstances where real-time performance is more important than strict Java and IEEE-standard compliance.

When VSCR[NJ] is 0 (default), operations are performed in Java mode. When VSCR[NJ] is 1, operations are carried out in the non-Java mode.

#### 3.2.1.1 Java Mode

Java compliance requires compliance with only a subset of the Java/IEEE/C9X standard. The Java subset helps simplify floating-point implementations, as follows:

- Reducing the number of operations that must be supported
- Eliminating exception status flags and traps
- Producing results corresponding to all disabled exceptions, thus eliminating enabling control flags
- Requiring only round-to-nearest rounding mode eliminates directed rounding modes and the associated rounding control flags.

Java compliance requires the following aspects of the IEEE standard:

- Supporting denorms as inputs and results (gradual underflow) for arithmetic operations
- Providing NaN results for invalid operations
- NaNs compare unordered with respect to everything, so that the result of any comparison of any NaN to any data type is always false.

In some implementations, floating-point operations in Java mode may have somewhat longer latency on normal operands and possibly much longer latency on denormalized operands than operations in non-Java mode. This means that in Java mode overall real-time response may be somewhat worse and deadline scheduling may be subject to much larger variance than non-Java mode.

### 3.2.1.2 Non-Java Mode

In the non-Java/non-IEEE/non-C9X mode ( $VSCR[NJ] = 1$ ), gradual underflow is not performed. Instead, any instruction that would have produced a denormalized result in Java mode substitutes a correctly signed zero ( $\pm 0.0$ ) as the final result. Also, denormalized input operands are flushed to the correctly signed zero ( $\pm 0.0$ ) before being used by the instruction.

The intent of this mode is to give programmers a way to assure optimum, data-insensitive, real-time response across implementations. Another way to improved response time would be to implement denormalized operations through software emulation.

It is architecturally permitted, but strongly discouraged, for an implementation to implement only non-Java mode. In such an implementation, the  $VSCR[NJ]$  does not respond to attempts to clear it and is always read back as a 1.

No other architecturally visible, implementation-specific deviations from this specification are permitted in either mode.

## 3.2.2 Floating-Point Infinities

Valid operations on infinities are processed according to the IEEE standard.

## 3.2.3 Floating-Point Rounding

All AltiVec floating-point arithmetic instructions use the IEEE default rounding mode, round-to-nearest. The IEEE directed rounding modes are not provided.

## 3.2.4 Floating-Point Exceptions

The following floating-point exceptions may occur during execution of AltiVec floating-point instructions.

- NaN operand exception
- Invalid operation exception
- Zero divide exception
- Log of zero exception
- Overflow exception
- Underflow exception

If an exception occurs, a result is placed into the corresponding target element as described in the following subsections. This result is the default result specified by Java, the IEEE standard, or C9X, as applicable. Recall that denormalized source values are treated as if they were zero when  $VSCR[NJ] = 1$ . The consequences regarding exceptions are as follows:

- Exceptions that can be caused by a zero source value can be caused by a denormalized source value when  $VSCR[NJ] = 1$ .
- Exceptions that can be caused by a nonzero source value cannot be caused by a denormalized source value when  $VSCR[NJ] = 1$ .

### 3.2.4.1 NaN Operand Exception

If the exponent of a floating-point number is 255 and the fraction is non-zero, then the value is a NaN. If the most significant bit of the fraction field of a NaN is zero, then the value is a signaling NaN (SNaN), otherwise it is a quiet NaN (QNaN). In all cases the sign of a NaN is irrelevant.

A NaN operand exception occurs when a source value for any of the following instructions is a NaN:

- An AltiVec instruction that would normally produce floating-point results
- Either of the two, Vector Convert to Unsigned Fixed-Point Word Saturate (**vctuxs**) or Vector Convert to Signed Fixed-Point Word Saturate (**vctxsxs**) instructions
- Any of the four vector floating-point compare instructions.

The following actions can be taken:

- If the AltiVec instruction would normally produce floating-point results, the corresponding result is a source NaN selected as follows. In all cases, if the selected source NaN is an SNaN, it is converted to the corresponding QNaN (by setting the high-order bit of the fraction field to 1 before being placed into the target element).

```

if the element in register vA is a NaN
    then the result is that NaN
else if the element in register vB is a NaN
    then the result is that NaN
else if the element in register vC is a NaN
    then the result is that NaN

```

- If the instruction is either of the two vector convert to fixed-point word instructions (**vctuxs**, **vctxsxs**), the corresponding result is 0x0000\_0000.  $VSCR[SAT]$  is not affected.

- If the instruction is Vector Compare Bounds Floating-Point (**vcmpbfp**[.]), the corresponding result is 0xC000\_0000.
- If the instruction is one of the other three vector floating-point compare instructions (**vcmpeqfp**[.], **vcmpfgefp**[.], **vcmpbfp**[.]), the corresponding result is 0x0000\_0000.

### 3.2.4.2 Invalid Operation Exception

An invalid operation exception occurs when a source value is invalid for the specified operation. The invalid operations are as follows:

- Magnitude subtraction of infinities
- Multiplication of infinity by zero
- Vector Reciprocal Square Root Estimate Float (**vrsqrtefp**) of a negative, nonzero number or -X
- Log base 2 estimate (**vlogefp**) of a negative, nonzero number or -X

The corresponding result is the QNaN 0x7FC0\_0000. This is the single-precision format analogy of the double precision format generated QNaN described in Chapter 3, “Operand Conventions,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 3.2.4.3 Zero Divide Exception

A zero divide exception occurs when a Vector Reciprocal Estimate Floating-Point (**vrefp**) or Vector Reciprocal Square Root Estimate Floating-Point (**vrsqrtefp**) instruction is executed with a source value of zero.

The corresponding result is infinity, where the sign is the sign of the source value, as follows:

- $1/+0.0 \rightarrow +\infty$
- $1/-0.0 \rightarrow -\infty$
- $1/(\sqrt{+0.0}) \rightarrow +\infty$
- $1/(\sqrt{-0.0}) \rightarrow -\infty$

### 3.2.4.4 Log of Zero Exception

A log of zero exception occurs when a Vector Log Base 2 Estimate Floating-Point instruction (**vlogefp**) is executed with a source value of zero. The corresponding result is infinity. The exception cases are as follows:

- **vlogefp**  $\log_2(\pm 0.0) \rightarrow -\infty$
- **vlogefp**  $\log_2(-x) \rightarrow \text{QNaN}$ , where  $x \neq 0$

### 3.2.4.5 Overflow Exception

An overflow exception happens when either of the following conditions occurs:

- For an AltiVec instruction that would normally produce floating-point results, the magnitude of what would have been the result if the exponent range were unbounded exceeds that of the largest finite single-precision number.
- For either of the two Vector Convert To Fixed-Point Word instructions (**vctuxs**, **vctxsx**), either a source value is an infinity or the product of a source value and 2 unsigned immediate value (UIMM) is a number too large to be represented in the target integer format.

The following actions can be taken:

- If the AltiVec instruction would normally produce floating-point results, the corresponding result is infinity, where the sign is the sign of the intermediate result.
- If the instruction is Vector Convert to Unsigned Fixed-Point Word Saturate (**vctuxs**), the corresponding result is 0xFFFF\_FFFF if the source value is a positive number or +X, and is 0x0000\_0000 if the source value is a negative number or -X. VSCR[SAT] is set.
- If the instruction is Vector Convert to Signed Fixed-Point Word Saturate (**vctfsx**), the corresponding result is 0x7FFF\_FFFF if the source value is a positive number or +X, and is 0x8000\_0000 if the source value is a negative number or -X. VSCR[SAT] is set.

### 3.2.4.6 Underflow Exception

Underflow exceptions occur only for AltiVec instructions that would normally produce floating-point results. Underflow is detected before rounding. Underflow occurs when a nonzero intermediate result, computed as though both the precision and the exponent range were unbounded, is less in magnitude than the smallest normalized single-precision number ( $2^{-126}$ ).

The following actions can be taken:

- If VSCR[NJ] = 0, the corresponding result is the value produced by denormalizing and rounding the intermediate result.
- If VSCR[NJ] = 1, the corresponding result is a zero, where the sign is the sign of the intermediate result.

### 3.2.5 Floating-Point NaNs

The AltiVec floating-point data format is compliant with the Java/IEEE/C9X single-precision format. A quantity in this format can represent a signed normalized number, a signed denormalized number, a signed zero, a signed infinity, a quiet not a number (QNaN), or a signaling NaN (SNaN).

### 3.2.5.1 NaN Precedence

Whenever only one source operand of an instruction that returns a floating-point result is a NaN, then that NaN is selected as the input NaN to the instruction. When more than one source operand is a NaN, the precedence order for selecting the NaN is first from **vA** then from **vB** and then from **vC**. If the selected NaN is an SNaN, it is processed as described in Section 3.2.5.2, “SNaN Arithmetic.” QNaN’s, are processed according to Section 3.2.5.3, “QNaN Arithmetic.”

### 3.2.5.2 SNaN Arithmetic

Whenever the input NaN to an instruction is an SNaN, a QNaN is delivered as the result, as specified by the IEEE standard when no trap occurs. The delivered QNaN is an exact copy of the original SNaN except that it is quieted; that is, the most-significant bit (msb) of the fraction is a one.

### 3.2.5.3 QNaN Arithmetic

Whenever the input NaN to an instruction is a QNaN, it is propagated as the result according to the IEEE standard. All information in the QNaN is preserved through all arithmetic operations.

### 3.2.5.4 NaN Conversion to Integer

All NaNs convert to zero on conversions to integer instructions such as **vctuxs** and **vctxsx**.

### 3.2.5.5 NaN Production

Whenever the result of an AltiVec operation is a NaN (for example, an invalid operation), the NaN produced is a QNaN with the sign bit = 0, exponent field = 255, msb of the fraction field = 1, and all other bits = 0.

# Chapter 4

## Addressing Modes and Instruction Set Summary

This chapter describes instructions and addressing modes defined by AltiVec Instruction Set Architecture (ISA) and according to the levels used by PowerPC architecture—user instruction set architecture (UISA) and virtual environment architecture (VEA). AltiVec instructions are primarily UISA; if otherwise, they are noted in the chapter. These instructions are divided into the following categories:



- Vector integer arithmetic instructions—These include arithmetic, logical, compare, rotate, and shift instructions, described in Section 4.2.1, “Vector Integer Instructions.”
- Vector floating-point arithmetic instructions—These include floating-point arithmetic instructions as well as a discussion on floating-point modes, described in Section 4.2.2, “Vector Floating-Point Instructions.”
- Vector load and store instructions—These include load and store instructions for vector registers, described in Section 4.2.3, “Load and Store Instructions.”
- Vector permutation and formatting instructions—These include pack, unpack, merge, splat, permute, select, and shift instructions, described in Section 4.2.5, “Vector Permutation and Formatting Instructions.”
- Processor control instructions—These instructions are used to read and write from the AltiVec Status and Control Register, described in Section 4.2.6, “Processor Control Instructions—UISA.”
- Memory control instructions—These instructions are used for managing caches (user level and supervisor level), described in Section 4.3.1, “Memory Control Instructions—VEA.”

This grouping of instructions does not necessarily indicate the execution unit that processes a particular instruction or group of instructions within a processor implementation.

AltiVec integer instructions operate on byte, half-word, and word operands. Floating-point instructions operate on single-precision operands. AltiVec ISA uses word-length instructions that are word-aligned. It provides for byte, half-word, and word operand fetches and stores between memory and the vector registers (VRs).

## Conventions

Arithmetic and logical instructions do not read or modify memory. To use the contents of a memory location in a computation for an arithmetic or logical instruction, the following steps are taken:

1. The memory contents must be loaded into a register with a load instruction.
2. The contents are then modified.
3. The modified contents are written to the target location using a store instruction.

## 4.1 Conventions

This section describes conventions used for the AltiVec instruction set. Descriptions of memory addressing, synchronization, and the AltiVec exception summary follow.

### 4.1.1 Execution Model

When used with PowerPC instructions, AltiVec instructions can be viewed as simply new PowerPC instructions that are freely intermixed with existing ones to provide additional functionality. Processors that implement the PowerPC architecture appear to execute instructions in program order. Some AltiVec implementations may not allow out-of-order execution and completion. Non-data dependent vector instructions may issue and execute while longer latency instructions issued previously are still in the execute stage. Register renaming avoids stalling dispatch on false dependencies and allows maximum register name reuse in heavily unrolled loops. The execution of a sequence of instructions will not be interrupted by exceptions since the unit does not report IEEE exceptions, but rather produces the default results as specified in the Java/IEEE/C9X standards. The execution of a sequence of instructions may be interrupted only by a vector load or store instruction; otherwise, AltiVec instructions do not generate any exceptions.

### 4.1.2 Computation Modes

AltiVec ISA supports the PowerPC ISA. The AltiVec ISA supports the 32-bit implementation of the PowerPC architecture in that all registers except FPRs and VRs are

**U** 32 bits long and the effective addresses are 32 bits long.

This chapter describes only the instructions defined for 32-bit implementations of the PowerPC architecture.

### 4.1.3 Classes of Instructions

AltiVec instructions follow the illegal instruction class defined by PowerPC architecture in the section, “Classes of Instructions,” in Chapter 4, “Addressing Modes and Instruction Set Summary,” of the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*. For AltiVec ISA, all unspecified encodings within the major opcode



(04) that are not defined are illegal PowerPC instructions. The only exclusion in defining an unspecified encoding is an unused bit in an immediate field or specifier field (///).

## 4.1.4 Memory Addressing

A program references memory using the effective (logical) address computed by the processor when it executes a load, store, or cache instruction, and when it fetches the next sequential instruction.

### 4.1.4.1 Memory Operands

Bytes in memory are numbered consecutively starting with zero. Each number is the address of the corresponding byte.

Memory operands may be bytes, half words, words, or quad words for AltiVec instructions. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Operand length is implicit for each instruction. AltiVec ISA supports both big-endian and little-endian byte ordering. The default byte and bit ordering is big-endian; see Section 3.1.2, “AltiVec Byte Ordering,” for more information.

The natural alignment boundary of an operand of a single-register memory access instruction is equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. For a detailed discussion about memory operands, see Section 3.1, “Data Organization in Memory.”

### 4.1.4.2 Effective Address Calculation

An effective address (EA) is the 32-bit sum computed by the processor when executing a memory access or when fetching the next sequential instruction. For a memory access instruction, if the sum of the EA and the operand length exceeds the maximum EA, the memory operand is considered to wrap around from the maximum EA through EA 0, as described in the Chapter 4, “Addressing Modes and Instruction Set Summary,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

A zero in the **rA** field indicates the absence of the corresponding address component. For the absent component, a value of zero is used for the address. This is shown in the instruction description as (**rA**|0).

In all implementations of processors that support the PowerPC architecture, the processor can modify the three low-order bits of the calculated effective address before accessing memory if the system is operating in little-endian mode. The double words of a quad word may be swapped as well. See Section 3.1.2, “AltiVec Byte Ordering,” for more information about little-endian mode.

AltiVec load and store operations use register indirect with index mode and boundary align to generate effective addresses. For further details see Section 4.2.3.2, “Load and Store Address Generation.”

## 4.2 AltiVec USA Instructions

AltiVec instructions can provide additional supporting instructions to PowerPC architecture. This section discusses the instructions defined in AltiVec user instruction set architecture (UISA).

### 4.2.1 Vector Integer Instructions

The following are categories for vector integer instructions:

- Arithmetic
- Compare
- Logical
- Rotate and shift

Integer instructions use the content of the vector registers (VRs) as source operands and place results into VRs as well. Setting the Rc bit of a vector compare instruction causes the PowerPC condition register (CR) to be updated.

AltiVec integer instructions treat source operands as signed integers unless the instruction is explicitly identified as performing an unsigned operation. For example, Vector Add Unsigned Word Modulo (**vadduwm**) and Vector Multiply Odd Unsigned Byte (**vmuloub**) instructions interpret both operands as unsigned integers.

#### 4.2.1.1 Saturation Detection

Most integer instructions have both signed and unsigned versions and many have both modulo (wrap-around) and saturating clamping modes. Saturation occurs whenever the result of a saturating instruction does not fit in the result field. Unsigned saturation clamps results to zero on underflow and to the maximum positive integer value ( $2^n-1$ , for example, 255 for byte fields) on overflow. Signed saturation clamps results to the smallest representable negative number ( $-2^{n-1}$ , for example, -128 for byte fields) on underflow, and to the largest representable positive number ( $2^{n-1}-1$ , for example, +127 for byte fields) on overflow. When a modulo instruction is used, the resultant number truncates overflow or underflow for the length (byte, half word, word, quad word) and type of operand (unsigned, signed). The AltiVec ISA provides a way to detect saturation and sets the SAT bit in the Vector Status and Control Register (VSCR[SAT]) in a saturating instruction.

Borderline cases that generate results equal to saturation values, for example unsigned  $0+0 \rightarrow 0$  and unsigned byte  $1+254 \rightarrow 255$ , are not considered saturation conditions and do not cause VSCR[SAT] to be set.

The VSCR[SAT] can be set by the following types of integer, floating-point, and formatting instructions:

- Move to VSCR (**mtvscr**)
- Vector add integer with saturation (**vaddubs**, **vadduhs**, **vadduws**, **vaddsubs**, **vaddshs**, **vaddsws**)
- Vector subtract integer with saturation (**vsububs**, **vsubuhs**, **vsubuws**, **vsubsubs**, **vsubshs**, **vsubsws**)
- Vector multiply-add integer with saturation (**vmhaddshs**, **vmhraddshs**)
- Vector multiply-sum with saturation (**vmsumuhs**, **vmsumshs**, **vsumsws**)
- Vector sum-across with saturation (**vsumsws**, **vsum2sws**, **vsum4sbs**, **vsum4shs**, **vsum4ubs**)
- Vector pack with saturation (**vpkuhus**, **vpkuwus**, **vpkshus**, **vpkswus**, **vpkshss**, **vpkswss**)
- Vector convert to fixed-point with saturation (**vctuxs**, **vctxsx**)

Note that only instructions that explicitly call for saturation can set VSCR[SAT]. Modulo integer instructions and floating-point arithmetic instructions never set VSCR[SAT]. For further details see Section 2.3.2, “Vector Status and Control Register (VSCR).”

#### 4.2.1.2 Vector Integer Arithmetic Instructions

Table 4-1 lists the integer arithmetic instructions for processors that implement the PowerPC architecture.

**Table 4-1. Vector Integer Arithmetic Instructions**

Name	Mnemonic	Syntax	Operation
Vector Add Unsigned Integer [b,h,w] Modulo	<b>vaddubm</b> <b>vadduhm</b> <b>vadduwm</b>	vD,vA,vB	<p>Places the sum (vA[unsigned integer elements]) + (vB[unsigned integer elements]) into vD[unsigned integer elements] using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add sixteen unsigned integers from vA to the corresponding sixteen unsigned integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add eight unsigned integers from vA to the corresponding eight unsigned integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add four unsigned integers from vA to the corresponding four unsigned integers from vB.</p> <p>Note: unsigned or signed integers can be used with these instructions.</p>
Vector Add Unsigned Integer [b,h,w] Saturate	<b>vaddubs</b> <b>vadduhs</b> <b>vadduws</b>	vD,vA,vB	<p>Place the sum (vA[unsigned integer elements]) + (vB[unsigned integer elements]) into vD[unsigned integer elements] using saturate clamping mode. Saturate clamping mode means if the resulting sum is <math>&gt;(2^n-1)</math> saturate to <math>(2^n-1)</math>, where <math>n = \mathbf{b,h,w}</math>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add sixteen unsigned integers from vA to the corresponding sixteen unsigned integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add eight unsigned integers from vA to the corresponding eight unsigned integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add four unsigned integers from vA to the corresponding four unsigned integers from vB.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Add Signed Integer[b,h,w] Saturate	<b>vaddsbs</b> <b>vaddshs</b> <b>vddsws</b>	vD,vA,vB	<p>Place the sum (vA[signed integer elements]) + (vB[signed integer elements]) into vD[signed integer elements] using saturate clamping mode. Saturate clamping mode means:</p> <p>if the sum is <math>&gt;(2^{n-1}-1)</math> saturate to <math>(2^{n-1}-1)</math> and if <math>&lt; (-2^{n-1})</math> saturate to <math>(-2^{n-1})</math>, where <math>n = \mathbf{b,h,w}</math>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = byte, add sixteen signed integers from vA to the corresponding sixteen signed integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add eight signed integers from vA to the corresponding eight signed integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add four signed integers from vA to the corresponding four signed integers from vB.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Add and Write Carry-out Unsigned Word	<b>vaddcuw</b>	vD,vA,vB	<p>Take the carry out of summing (vA) + (vB) and place it into vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 2 bytes, add four unsigned integers from vA to the corresponding four unsigned integers from vB and the resulting carry outs are correspondingly placed in vD.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Subtract Unsigned Integer Modulo [b,h,w]	<b>vsububm</b> <b>vsubuhm</b> <b>vsubuwm</b>	<b>vD,vA,vB</b>	Place the unsigned integer sum (vA) - (vB) into vD using modulo arithmetic.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, subtract sixteen unsigned integers in vB from the corresponding sixteen unsigned integers in vA.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, subtract eight unsigned integers in vB from the corresponding eight unsigned integers in vA.  For <b>w</b> , word, integer length = 32 bits = 4 bytes, subtract four unsigned integers in vB from the corresponding four unsigned integers in vA.  Note that unsigned or signed integers can be used with these instructions.
Vector Subtract Unsigned Integer Saturate [b,h,w]	<b>vsububs</b> <b>vsubuhs</b> <b>vsubuws</b>	<b>vD,vA,vB</b>	Place the unsigned integer sum vA - vB into vD using saturate clamping mode, that is, if the sum < 0, it saturates to 0 corresponding to <b>b,h,w</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, subtract sixteen unsigned integers in vB from the corresponding sixteen unsigned integers in vA.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, subtract eight unsigned integers in vB from the corresponding eight unsigned integers in vA.  For <b>w</b> , word, integer length = 32 bits = 4 bytes, subtract four unsigned integers in vB from the corresponding four unsigned integers in vA.  If the result saturates, VSCR[SAT] is set.
Vector Subtract Signed Integer Saturate [b,h,w]	<b>vsubsbs</b> <b>vsubshs</b> <b>vsubsws</b>	<b>vD,vA,vB</b>	Place the signed integer sum (vA) - (vB) into vD using saturate clamping mode. Saturate clamping mode means:  if the sum is >(2 <sup>n-1</sup> -1) saturate to (2 <sup>n-1</sup> -1) and if < (- 2 <sup>n-1</sup> ) saturate to (-2 <sup>n-1</sup> ), where n= <b>b,h,w</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, subtract sixteen signed integers in vB from the corresponding sixteen signed integers in vA.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, subtract eight signed integers in vB from the corresponding eight signed integers in vA.  For <b>w</b> , word, integer length = 32 bits = 4 bytes, subtract four signed integers in vB from the corresponding four signed integers in vA.
Vector Subtract and Write Carry-out Unsigned Word	<b>vsubcuw</b>	<b>vD,vA,vB</b>	Take the carry out of the sum (vA) - (vB) and place it into vD.  For <b>w</b> , word, integer length = 32 bits = 2 bytes, subtract four unsigned integers in vB from the corresponding four unsigned integers in vA and place the resulting carry outs into vD.

**Table 4-1. Vector Integer Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply Odd Unsigned Integer [b,h] Modulo	<b>vmuloub</b> <b>vmulouh</b>	<b>vD,vA,vB</b>	Place the unsigned integer products of ( <b>vA</b> ) * ( <b>vB</b> ) into <b>vD</b> using modulo arithmetic mode.  For <b>b</b> , byte, integer length = 8 bits =1 byte, multiply 8 odd-numbered unsigned integer byte elements from <b>vA</b> to the corresponding 8 odd-numbered unsigned integer byte elements from <b>vB</b> resulting in eight unsigned integer half-word products in <b>vD</b> .  For <b>h</b> , half word, integer length =16 bits = 2 bytes, multiply 4 odd-numbered unsigned integer half word elements from <b>vA</b> to the corresponding 4 odd numbered unsigned integer half-word elements from <b>vB</b> resulting in four unsigned integer word products in <b>vD</b> .
Vector Multiply Odd Signed Integer [b,h] Modulo	<b>vmulosb</b> <b>vmulosh</b>	<b>vD,vA,vB</b>	Place the signed integer product of ( <b>vA</b> ) * ( <b>vB</b> ) into <b>vD</b> using modulo arithmetic mode.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, multiply 8 odd-numbered signed integer byte elements from <b>vA</b> to 8 odd-numbered signed integer byte elements from <b>vB</b> resulting in eight signed integer half-word products in <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, multiply 4 odd-numbered signed integer half word elements from <b>vA</b> to 4 odd-numbered signed integer half word elements from <b>vB</b> resulting in four signed integer word products in <b>vD</b> .
Vector Multiply Even Unsigned Integer [b,h] Modulo	<b>vmuleub</b> <b>vmuleuh</b>	<b>vD,vA,vB</b>	Place the unsigned integer products of ( <b>vA</b> ) * ( <b>vB</b> ) into <b>vD</b> using modulo arithmetic mode.  For <b>b</b> , byte, integer length = 8 bits =1 byte, multiply 8 even-numbered unsigned integer byte elements from <b>vA</b> to 8 even-numbered unsigned integer byte elements from <b>vB</b> resulting in eight unsigned integer half-word products in <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, multiply 4 even-numbered unsigned integer half-word elements from <b>vA</b> to 4 even numbered unsigned integer half- word elements from <b>vB</b> resulting in four unsigned integer word products in <b>vD</b>
Vector Multiply Even Signed Integer [b,h] Modulo	<b>vmulesb</b> <b>vmulesh</b>	<b>vD,vA,vB</b>	Place the signed integer product of ( <b>vA</b> ) * ( <b>vB</b> ) into <b>vD</b> using modulo arithmetic mode.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, multiply 8 even-numbered signed integer byte elements from <b>vA</b> to 8 even-numbered signed integer byte elements from <b>vB</b> resulting in eight signed integer half-word products in <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, multiply 4 even-numbered signed integer half-word elements from <b>vA</b> to 4 even-numbered signed integer half-word elements from <b>vB</b> resulting in four signed integer word products in <b>vD</b> .

**Table 4-1. Vector Integer Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply-High and Add Signed Half-Word Saturate	<b>vmhaddshs</b>	<b>vD,vA,vB, vC</b>	<p>The 17 most significant bits (msb's) of the product of (vA) * (vB) adds to sign-extended vC and places the result into vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the eight signed half words from vA with the corresponding eight signed half words from vB to produce a 32-bit intermediate product and then take the 17 msb's (bits 0–16) of the 8 intermediate products and add them to the 8 sign-extended half words in vC, place the 8 half-word saturated results in vD. If the intermediate product is as follows:</p> <p>&gt; <math>(2^{15}-1)</math> saturate to <math>(2^{15}-1)</math> and if                      &lt; <math>-2^{15}</math> saturate to <math>-2^{15}</math>.</p> <p>If the results saturates, VSCR[SAT] is set.</p>
Vector Multiply-High Round and Add Signed Half-Word Saturate	<b>vmhraddshs</b>	<b>vD,vA,vB,vC</b>	<p>Add the rounded product of (vA) * (vB) to sign-extended vC and place the result into vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the eight signed integers from vA to the corresponding eight signed integers from vB and then round the 8 immediate products by adding the value 0x0000_4000 to it. Then add the most significant bits (msb), bits 0–16, of the 8 rounded immediate products to the 8 sign-extended values in vC and place the eight signed half-word saturated results into vD. If the intermediate product is:</p> <p>&gt; <math>(2^{15}-1)</math> saturate to <math>(2^{15}-1)</math> or if                      &lt; <math>-2^{15}</math> saturate to <math>-2^{15}</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Multiply-Low and Add Unsigned Half-Word Modulo	<b>vmladduhm</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to zero-extended vC and place into vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the eight signed integers from vA to the corresponding eight signed integers from vB to produce a 32-bit intermediate product. The 16-bit value in vC is zero-extended to 32 bits and added to the intermediate product and the lower 16 bits of the sum (bit 16–31) is placed in vD.</p> <p>Note that unsigned or signed integers can be used with these instructions.</p>
Vector Multiply-Sum Unsigned Integer [b,h] Modulo	<b>vmsumubm</b> <b>vmsumuhm</b>	<b>vD,vA,vB,vC</b>	<p>The product of (vA) * (vB) is added to zero-extended vC and placed into vD using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply four unsigned integer bytes from a word element in vA by the corresponding four unsigned integer bytes in a word element in vB and the sum of these products are added to the zero-extended unsigned integer word element in vC and then placed the unsigned integer word result into vD, following this process for each 4-word element in vA and vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 unsigned integer half words from a word element in vA by the corresponding 2 unsigned integer half words in a word element in vB and the sum of these products are added to zero-extended unsigned integer word element in vC and then place the unsigned integer word result into vD, following this process for each 4 word element in vA and vB.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply-Sum Signed Half-Word Saturate	<b>vmsumshs</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to vC and place the result into vD using saturate clamping mode.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 signed integer half words from a word element in vA by the corresponding 2 signed integer half words in a word element in vB. Add the sum of these products to the signed integer word element in vC and then place the signed integer word result into vD, (following this process for each 4-word element in vA and vB). If the intermediate result is &gt; (2<sup>31</sup>-1), saturate to (2<sup>31</sup>-1) and if the result is &lt; -2<sup>31</sup>, saturate to -2<sup>31</sup>. If the result saturates, VSCR[SAT] is set.</p>
Vector Multiply-Sum Unsigned Half-Word Saturate	<b>vmsumuhs</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to zero-extended vC and place the result into vD using saturate clamping mode.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 unsigned integer half words from a word element in vA by the corresponding 2 unsigned integer half words in a word element in vB. Add the sum of these products to the zero-extended unsigned integer word element in vC and then place the unsigned integer word result into vD, (following this process for each 4-word element in vA and vB). If the intermediate result is &gt; (2<sup>32</sup>-1) saturate to (2<sup>32</sup>-1). If the result saturates, VSCR[SAT] is set.</p>
Vector Multiply-Sum Mixed Sign Byte Modulo	<b>vmsummbm</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to vC and place into vD using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply four signed integer bytes from a word element in vA by the corresponding four unsigned integer bytes from a word element in vB. Add the sum of these four signed products to the signed integer word element in vC and then place the signed integer word result into vD, following this process for each 4-word element in vA and vB.</p>
Vector Multiply-Sum Signed Half-Word Modulo	<b>vmsumshm</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to vC and place into vD using modulo arithmetic.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 signed integer half words from a word element in vA by the corresponding 2 signed integer half words in a word element in vB. Add the sum of these 2 products to the signed integer word element in vC and then place the signed integer word result into vD, following this process for each 4-word element in vA and vB.</p>
Vector Sum Across Signed Word Saturate	<b>vsumsws</b>	<b>vD,vA,vB</b>	<p>Place the sum of signed word elements in vA and the word in vB[96-127] into vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add the sum of the four signed integer word elements in vA to the word element in vB[96-127]. If the intermediate product is &gt; (2<sup>31</sup>-1) saturate to (2<sup>31</sup>-1) and if &lt; -2<sup>31</sup> saturate to -2<sup>31</sup>. Place the signed integer result in vD[96-127], vD[0-95] are cleared.</p>



Table 4-1. Vector Integer Arithmetic Instructions (continued)

Name	Mnemonic	Syntax	Operation
Vector Sum Across Partial (1/2) Signed Word Saturate	<b>vsum2sws</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[word 0 + word 1] + <b>vB</b>[word 1] and place in <b>vD</b>[word 1]. Repeat only add <b>vA</b>[word 2 + word 3] + <b>vB</b>[word 3] and place in <b>vD</b>[word 3].</p> <p>word 0 = Bits 0–31                      word 1 = Bits 32-63                      word 2 = Bits 64-95                      word 3 = Bits 96-127,</p> <p>Figure1-2 shows a picture of what the word elements would look like in a vector register.</p> <p>Add the sum of word 0 and word 1 of <b>vA</b> to word 1 of <b>vB</b> using saturate clamping mode and place the result is into word 1 of <b>vD</b>. Then add the sum of word 2 and word 3 of (<b>vA</b>) to word 3 of <b>vB</b> using saturate clamping mode and place those results into word 3 in <b>vD</b>. If the intermediate result for either calculation is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>If the result saturates, <b>VSCR[SAT]</b> is set.</p>
Vector Sum Across Partial (1/4) Unsigned Byte Saturate	<b>vsum4ubs</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[4 byte elements sum to a word] and <b>vB</b>[word element] then place in <b>vD</b>[word element] using saturate clamping mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, for each word element in <b>vB</b>, add the sum of four unsigned bytes in the word in <b>vA</b> to the unsigned word element in <b>vB</b> and then place the results into the corresponding unsigned word element in <b>vD</b>. If the intermediate result for is <math>&gt; (2^{32}-1)</math> it saturates to <math>(2^{32}-1)</math>.</p> <p>If the result saturates, <b>VSCR[SAT]</b> is set.</p>
Vector Sum Across Partial (1/4) Signed Integer Saturate	<b>vsum4sbs</b> <b>vsum4shs</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[sum of signed integer elements in word] and <b>vB</b>[word element] then place in <b>vD</b>[word element] using saturate clamping mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, for each word element in <b>vB</b>, add the sum of four signed bytes in the word in <b>vA</b> to the signed word element in <b>vB</b> and then place the results into the corresponding signed word element in <b>vD</b>. If the intermediate result is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, for each word element in <b>vB</b>, add the sum of 2 signed half words in the word in <b>vA</b> to the signed word element in <b>vB</b> and then place the results into the corresponding signed word element in <b>vD</b>. If the intermediate result is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>If the result saturates, <b>VSCR[SAT]</b> is set.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Average Unsigned Integer [b,h,w]	<b>vavgub</b> <b>vavguh</b> <b>vavguw</b>	<b>vD,vA,vB</b>	<p>Add the sum of (vA[unsigned integer elements]+ vB[unsigned integer elements]) +1 and place into vD using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add sixteen unsigned integers from vA to sixteen unsigned integers from vB and then add 1 to the sums and place the high order result in vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add eight unsigned integers from vA to eight unsigned integers from vB and then add 1 to the sums and place the high order result in vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add four unsigned integers from vA to four unsigned integers from vB and then add 1 to the sums and place the high order result in vD.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Average Signed Integer [b,h,w]	<b>vavgsb</b> <b>vavgsh</b> <b>vavgsw</b>	<b>vD,vA,vB</b>	<p>Add the sum of (vA[signed integer elements]+ vB[signed integer elements]) +1 and place into vD using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add sixteen signed integers from vA to sixteen signed integers from vB and then add 1 to the sums and place the high order result in vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add eight signed integers from vA to eight signed integers from vB and then add 1 to the sums and place the high order result in vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add four signed integers from vA to four signed integers from vB and then add 1 to the sums and place the high order result in vD.</p>
Vector Maximum Unsigned Integer [b,h,w]	<b>vmaxub</b> <b>vmaxuh</b> <b>vmaxuw</b>	<b>vD,vA,vB</b>	<p>Compare the maximum of vA and vB unsigned integers for each integer value and which ever value is larger, place that unsigned integer value into vD</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare sixteen unsigned integers from vA with sixteen unsigned integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare eight unsigned integers from vA with eight unsigned integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare four unsigned integers from vA with four unsigned integers from vB.</p>
Vector Maximum Signed Integer [b,h,w]	<b>vmaxsb</b> <b>vmaxsh</b> <b>vmaxsw</b>	<b>vD,vA,vB</b>	<p>Compare the maximum of vA and vB signed integers for each integer value and which ever value is larger, place that signed integer value into vD</p> <p>For <b>b</b>, byte, integer length = 8 bits =1 byte, compare sixteen signed integers from vA with sixteen signed integers from vB.</p> <p>For <b>h</b>, half word, integer length =16 bits = 2 bytes, compare eight signed integers from vA with eight signed integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare four signed integers from vA with four signed integers from vB.</p>

Table 4-1. Vector Integer Arithmetic Instructions (continued)

Name	Mnemonic	Syntax	Operation
Vector Minimum Unsigned Integer [b,h,w]	<b>vminub</b> <b>vminuh</b> <b>vminuw</b>	<b>vD,vA,vB</b>	Compare the minimum of <b>vA</b> and <b>vB</b> unsigned integers for each integer value and which ever value is smaller, place that unsigned integer value into <b>vD</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, compare sixteen unsigned integers from <b>vA</b> with sixteen unsigned integers from <b>vB</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, compare eight unsigned integers from <b>vA</b> with eight unsigned integers from <b>vB</b> .  For <b>w</b> , word, integer length = 32 bits = 4 bytes, compare four unsigned integers from <b>vA</b> with four unsigned integers from <b>vB</b> .
Vector Minimum Signed Integer [b,h,w]	<b>vminsb</b> <b>vminsh</b> <b>vminsw</b>	<b>vD,vA,vB</b>	Compare the minimum of <b>vA</b> and <b>vB</b> signed integers for each integer value and which ever value is smaller, place that signed integer value into <b>vD</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, compare sixteen signed integers from <b>vA</b> with sixteen signed integers from <b>vB</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, compare eight signed integers from <b>vA</b> with eight signed integers from <b>vB</b> .  For <b>w</b> , word, integer length = 32 bits = 4 bytes, compare four signed integers from <b>vA</b> with four signed integers from <b>vB</b> .

### 4.2.1.3 Vector Integer Compare Instructions

The vector integer compare instructions algebraically or logically compare the contents of the elements in vector register **vA** with the contents of the elements in **vB**. Each compare result vector is comprised of TRUE (0xFF, 0xFFFF, 0xFFFFFFFF) or FALSE (0x00, 0x0000, 0x00000000) elements of the size specified by the compare source operand element (byte, half word, or word). The result vector can be directed to any vector register and can be manipulated with any of the instructions as normal data, for example, combining condition results. Vector compares provide equal-to and greater-than predicates. Others are synthesized from these by logically combining or inverting result vectors.

If the record bit (**Rc**) is set in the integer compare instructions (shown in Table 4-3), it can optionally set the **CR6** field of the PowerPC condition register. If **Rc = 1** in the vector integer compare instruction, then **CR6** reflects the result of the comparison, as shown in Table 4-2.

Table 4-2. CR6 Field Bit Settings for Vector Integer Compare Instructions

CR Bit	CR6 Bit	Vector Compare
24	0	1 Relation is true for all element pairs (that is, <b>vD</b> is set to all ones).
25	1	0
26	2	1 Relation is false for all element pairs (that is, register <b>vD</b> is cleared).
27	3	0

Table 4-3 summarizes the vector integer compare instructions.

**Table 4-3. Vector Integer Compare Instructions**

Name	Mnemonic	Syntax	Operation
Vector Compare Greater than Unsigned Integer [b,h,w]	vcmpgtub[.] vcmpgtuh[.] vcmpgtuw[.]	vD,vA,vB	<p>Compare the value in vA with the value in vB, treating the operands as unsigned integers. Place the result of the comparison into the vD field specified by operand vD.</p> <p>If vA &gt; vB then vD = 1's; otherwise vD = 0's.</p> <p>If the record bit (Rc) is set in the vector compare instruction, then vD == 1's, (all elements true) then CR6[0] is set vD == 0's, (all elements false) then CR6[2] is set.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare sixteen unsigned integers from vA to sixteen unsigned integers from vB and place the results in the corresponding 16 elements in vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare eight unsigned integers from vA to eight unsigned integers from vB and place the results in the corresponding 8 elements in vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare four unsigned integers from vA to four unsigned integers from vB and place the results in the corresponding 4 elements in vD.</p>

Table 4-3. Vector Integer Compare Instructions (continued)

Name	Mnemonic	Syntax	Operation
Vector Compare Greater Than Signed Integer [b,h,w]	vcmpgtsb[.] vcmpgtsh[.] vcmpgtsw[.]	vD,vA,vB	<p>Compare the value in vA with the value in vB, treating the operands as signed integers. Place the result of the comparison into the vD field specified by operand vD.</p> <p>If vA &gt; vB then vD = 1's; otherwise vD = 0's</p> <p>If the record bit (Rc) is set in the vector compare instruction, then vD == 1's, (all elements true) then CR6[0] is set vD == 0's, (all elements false) then CR6[2] is set.</p> <p>For b, byte, integer length = 8 bits = 1 byte, compare sixteen signed integers from vA to sixteen signed integers from vB and place the results in the 16 corresponding elements in vD.</p> <p>For h, half word, integer length = 16 bits = 2 bytes, compare eight signed integers from vA to eight signed integers from vB and place the results in the 8 corresponding elements in vD.</p> <p>For w, word, integer length = 32 bits = 4 bytes, compare four signed integers from vA to four signed integers from vB and place the results in the 4 corresponding elements in vD.</p>
Vector Compare Equal To Unsigned Integer [b,h,w]	vcmpequb[.] vcmpequh[.] vcmpequw[.]	vD,vA,vB	<p>Compare the value in vA with the value in vB, treating the operands as unsigned integers. Place the result of the comparison into the vD field specified by operand vD.</p> <p>If vA = vB then vD = 1's; otherwise vD = 0's.</p> <p>If the record bit (Rc) is set in the vector compare instruction then vD == 1's, (all elements true) then CR6[0] is set vD == 0's, (all elements false) then CR6[2] is set.</p> <p>For b, byte, integer length = 8 bits = 1 byte, compare sixteen unsigned integers from vA to sixteen unsigned integers from vB and place the results in the corresponding 16 elements in vD.</p> <p>For h, half word, integer length = 16 bits = 2 bytes, compare eight unsigned integers from vA to eight unsigned integers from vB and place the results in the corresponding 8 elements in vD.</p> <p>For w, word, integer length = 32 bits = 4 bytes, compare four unsigned integers from vA to four unsigned integers from vB and place the results in the corresponding 4 elements in vD.</p> <p>Note: vcmpequb[.], vcmpequh[.], and vcmpequw[.] can use both unsigned and signed integers.</p>

#### 4.2.1.4 Vector Integer Logical Instructions

The vector integer logical instructions shown in Table 4-4 perform bit-parallel operations on the operands.

**Table 4-4. Vector Integer Logical Instructions**

Name	Mnemonic	Syntax	Operation
Vector Logical AND	vand	vD,vA,vB	AND the contents of vA with vB and place the result into vD.
Vector Logical OR	vor	vD,vA,vB	OR the contents of vA with vB and place the result into vD.
Vector Logical XOR	vxor	vD,vA,vB	XOR the contents of vA with vB and place the result into vD.
Vector Logical AND with Complement	vandc	vD,vA,vB	AND the contents of vA with the complement of vB and place the result into vD.
Vector Logical NOR	vnor	vD,vA,vB	NOR the contents of vA a with vB and place the result into vD.

### 4.2.1.5 Vector Integer Rotate and Shift Instructions

The vector integer rotate instructions are summarized in Table 4-5.

**Table 4-5. Vector Integer Rotate Instructions**

Name	Mnemonic	Syntax	Operation
Vector Rotate Left Integer [b,h,w]	vrlb vrlh vrlw	vD,vA,vB	<p>Rotate each element in vA left by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in vB. Place the result into the corresponding element of vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from vA with 16 integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from vA with 8 integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from vA with 4 integers from vB.</p>

The vector integer shift instructions are summarized in Table 4-6.

Table 4-6. Vector Integer Shift Instructions

Name	Mnemonic	Syntax	Operation
Vector Shift Left Integer [b,h,w]	vslb vslh vslw	vD,vA,vB	<p>Shift each element in vA left by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in vB. If bits are shifted out of bit 0 of the element they are lost. Supply zeros to the vacated bits on the right. Place the result into the corresponding element of vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from vA with 16 integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from vA with 8 integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from vA with 4 integers from vB.</p>
Vector Shift Right Integer [b,h,w]	vsrb vsrh vsrw	vD,vA,vB	<p>Shift each element in vA right by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in vB. If bits are shifted out of bit n-1 of the element they are lost. Supply zeros to the vacated bits on the left. Place the result into the corresponding element of vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from vA with 16 integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from vA with 8 integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from vA with 4 integers from vB.</p>
Vector Shift Right Algebraic Integer [b,h,w]	vsrab vsrah vsraw	vD,vA,vB	<p>Shift each element in vA right by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in vB. If bits are shifted out of bit n-1 of the element they are lost. Replicate bit 0 of the element to fill the vacated bits on the left. Place the result into the corresponding element of vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from vA with 16 integers from vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from vA with 8 integers from vB.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from vA with 4 integers from vB.</p>

### 4.2.2 Vector Floating-Point Instructions

This section describes the vector floating-point instructions, which include the following:

- Arithmetic
- Rounding and conversion
- Compare
- Estimate

The AltiVec floating-point data format complies with the ANSI/IEEE-754 standard. A quantity in this format represents a signed normalized number, a signed denormalized number, a signed zero, a signed infinity, a quiet not a number (QNaN), or a signalling NaN (SNaN). Operations perform to a Java/IEEE/C9X-compliant subset of the IEEE standard,

for further details on the Java or Non-Java mode see Section 3.2.1, “Floating-Point Modes.” AltiVec ISA does not report IEEE exceptions but rather produces default results as specified by the Java/IEEE/C9X Standard. For further details on exceptions, see Section 3.2.4, “Floating-Point Exceptions.”

### 4.2.2.1 Floating-Point Division and Square-Root

AltiVec instructions do not have division or square-root instructions. AltiVec ISA implements Vector Reciprocal Estimate Floating-Point (**vrefp**) and Vector Reciprocal-Square-Root Estimate Floating-Point (**vrsqrtefp**) instructions along with a Vector Negative Multiply-Subtract Floating-Point (**vnmsubfp**) instruction assisting in the Newton-Raphson refinement of the estimates. To accomplish division, simply multiply the dividend ( $x/y = x * 1/y$ ) and square-root by multiplying the original number ( $\sqrt{x} = x * 1/\sqrt{x}$ ). In this way, AltiVec ISA provides inexpensive divides and square-roots that are fully pipelined, sub-operation scheduled, and faster even than many hardware dividers. Software methods are available to further refine these to correct IEEE results.

#### 4.2.2.1.1 Floating-Point Division

The Newton-Raphson refinement step for the reciprocal  $1/B$  looks like this:

$$y1 = y0 + y0*(1 - B*y0), \text{ where } y0 = \text{recip\_est}(B)$$

This is implemented in the AltiVec ISA as follows:

```
y0 = vrefp(B)
t = vnmsubfp(y0,B,1)
y1 = vmaddfp(y0,t,y0)
```

This produces a result accurate to almost 24 bits of precision, except where B is a sufficiently small denormalized number that **vrefp** generates an infinity that, if important, must be explicitly guarded against.

To get a correctly rounded IEEE quotient from the above result, a second Newton-Raphson iteration is performed to get a correctly rounded reciprocal (y2) to the required 24 bits of precision, then the residual.

$$R = A - B*Q$$

is computed with **vnmsubfp** (where A is the dividend, B the divisor, and Q an approximation of the quotient from A\*y2). The correctly rounded quotient can then be obtained.

$$Q' = Q + R*y2$$

The additional accuracy provided by the fused nature of the AltiVec instruction multiply-add is essential to producing the correctly rounded quotient by this method.



The second Newton-Raphson iteration may ultimately not be needed but more work must be done to show that the absolute error after the first refinement step would always be less than 1 ulp, which is a requirement of this method.

#### 4.2.2.1.2 Floating-Point Square-Root

The Newton-Raphson refinement step for reciprocal square root looks like the following:

$$y1 = y0 + 0.5*y0*(1 - B*y0*y0), \quad \text{where } y0 = \text{recip\_sqrt\_est}(B)$$

That can be implemented as follows:

```

y0 = vrsqrtefp(B)
t0 = vmaddfp(y0,y0,0.0)
t1 = vmaddfp(y0,0.5,0.0)
t0 = vnmsubfp(B,t0,1)
y1 = vmaddfp(t0,t1,y0)
    
```

Various methods can further refine a correctly rounded IEEE result, all more elaborate than the simple residual correction for division, and therefore are not presented here, but most of which also benefit from the negative multiply-subtract instruction.

#### 4.2.2.2 Floating-Point Arithmetic Instructions

The floating-point arithmetic instructions are summarized in Table 4-7.

**Table 4-7. Floating-Point Arithmetic Instructions**

Name	Mnemonic	Syntax	Operation
Vector Add Floating-Point	vaddfp	vD,vA, vB	Add the 4-word (32-bit) floating-point elements in vA to the 4-word (32-bit) floating-point elements in vB. Round the four intermediate results to the nearest single-precision number and placed into vD.
Vector Subtract Floating-Point	vsubfp	vD,vA, vB	The 4-word (32-bit) floating-point values in vB are subtracted from the 4 32-bit values in vA. The four intermediate results are rounded to the nearest single-precision floating-point and placed into vD.

**Table 4-7. Floating-Point Arithmetic Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Vector Maximum Floating-Point	vmaxfp	vD,vA, vB	Compare each of the 4 single-precision word elements in vA to the corresponding 4 single-precision word elements in vB and place the larger value within each pair into the corresponding word element in vD. vmaxfp is sensitive to the sign of 0.0. When both operands are ±0.0: $\max(+0.0, \pm 0.0) = \max(\pm 0.0, +0.0) \Rightarrow +0.0$ $\max(-0.0, -0.0) \Rightarrow -0.0$ $\max(\text{NaN}, x) \Rightarrow \text{QNaN}$ , where x = any value
Vector Minimum Floating-Point	vminfp	vD,vA, vB	Compare each of the 4 single-precision word elements in vA to the corresponding 4 single-precision word elements in vB For each of the four elements, place the smaller value within each pair into vD. vminfp is sensitive to the sign of 0.0. When both operands are ±0.0: $\min(-0.0, \pm 0.0) = \min(\pm 0.0, -0.0) \Rightarrow -0.0$ $\min(+0.0, +0.0) \Rightarrow +0.0$ $\min(\text{NaN}, x) \Rightarrow \text{QNaN}$ where x = any value

### 4.2.2.3 Floating-Point Multiply-Add Instructions

Vector multiply-add instructions are critically important to performance because multiply followed by a data-dependent addition is the most common idiom in DSP algorithms. In most implementations, floating-point multiply-add instructions perform with the same latency as either a multiply or add alone, thus doubling performance in comparing to the otherwise serial multiply and adds. This will make performance twice as fast as using separate multiply and add instructions.

AltiVec floating-point multiply-adds instructions fuse (a multiply-add fuse implies that the full product participates in the add operation without rounding; only the final result rounds). This not only simplifies the implementation and reduces latency (by eliminating the intermediate rounding) but also increases the accuracy compared to separate multiply and adds.

Be careful as Java-compliant programs can not use multiply-add instructions fused directly because Java requires both the product and sum to round separately. Thus to achieve strict Java compliance, perform the multiply and add with separate instructions.

To realize multiply in AltiVec ISA use multiply-add instructions with a zero addend (for example, **vmaddfp** vD,vA,vC,vB where (vB = 0.0)).

Note that to use multiply-add instructions to perform an IEEE- or Java-compliant multiply, the addend must be -0.0. This is necessary to ensure that the sign of a zero result is correct when the product is either +0.0 or -0.0 ( $+0.0 + -0.0 \Rightarrow +0.0$ , and  $-0.0 + -0.0 \Rightarrow -0.0$ ). When the sign of a resulting 0.0 is not important, then use +0.0 as the addend that may, in some

cases, avoiding the need for a second register to hold a -0.0 in addition to the integer 0/floating-point +0.0 that may already be available.

The floating-point multiply-add instructions are summarized in Table 4-8.

**Table 4-8. Floating-Point Multiply-Add Instructions**

Name	Mnemonic	Syntax	Operation
Vector Multiply-Add Floating-Point	vmaddfp	vD,vA,vC,vB	Multiply the four word floating-point elements in vA by the corresponding four word elements in vC. Add the four word elements in vB to the four intermediate products. Round the results to the nearest single-precision numbers and place the corresponding word elements into vD.
Vector Negative Multiply-Subtract Floating-Point	vnmsubfp	vD,vA,vC,vB	Multiply the four word floating-point elements in vA by the corresponding four word elements in vC. Subtract the four word floating-point elements in vB from the four intermediate products and invert the sign of the difference. Round the results to the nearest single-precision numbers and place the corresponding word elements into vD.

#### 4.2.2.4 Floating-Point Rounding and Conversion Instructions

All AltiVec floating-point arithmetic instructions use the IEEE default rounding mode, round-to-nearest. AltiVec ISA does not provide the IEEE directed rounding modes.

AltiVec ISA provides separate instructions for converting floating-point numbers to integral floating-point values for all IEEE rounding modes as follows:

- Round-to-nearest (**vrfin**) (round)
- Round-toward-zero (**vrfiz**) (truncate)
- Round-toward-minus-infinity (**vrfim**) (floor)
- Round-toward-positive-infinity (**vrfip**) (ceiling).

Floating-point conversions to integers (**vctuxs**, **vctxsx**) use round-toward-zero (truncate). The floating-point rounding instructions are described in Table 4-9.

**Table 4-9. Floating-Point Rounding and Conversion Instructions**

Name	Mnemonic	Syntax	Operation
Vector Round to Floating-Point Integer Nearest	vrfin	vD,vB	Round to the nearest the four word floating-point elements in vB and place the four corresponding word elements into vD.
Vector Round to Floating-Point Integer toward Zero	vrfiz	vD,vB	Round towards zero the four word floating-point elements in vB and place the four corresponding word elements into vD.
Vector Round to Floating-Point Integer toward Positive Infinity	vrfip	vD,vB	Round towards +Infinity the four word floating-point elements in vB and place the four corresponding word elements into vD.

Table 4-9. Floating-Point Rounding and Conversion Instructions (continued)

Name	Mnemonic	Syntax	Operation
Vector Round to Floating-Point Integer toward Minus Infinity	vrfim	vD,vB	Round towards -Infinity the four word floating-point elements in vB and place the four corresponding word elements into vD.
Vector Convert from Unsigned Fixed-Point Word	vcfux	vD,vB, UIMM	Convert each of the four unsigned fixed-point integer word elements in vB to the nearest single-precision value. Divide the result by $2^{UIMM}$ and place into the corresponding word element of vD.
Vector Convert from Signed Fixed-Point Word	vcfsx	vD,vB, UIMM	Convert each signed fixed-point integer word element in vB to the nearest single-precision value. Divide the result by $2^{UIMM}$ and place into the corresponding word element of vD.
Vector Convert to Unsigned Fixed-Point Word Saturate	vctuxs	vD,vB, UIMM	Multiply each of the four single-precision word elements in vB by $2^{UIMM}$ . The products are converted to unsigned fixed-point integers using the Round toward Zero mode. If the intermediate results are $> 2^{32}-1$ saturate to $2^{32}-1$ and if it is $< 0$ saturate to 0. Place the unsigned integer results into the corresponding word elements of vD.
Vector Convert to Signed Fixed-Point Word Saturate	vctsxs	vD,vB, UIMM	Multiply each of the four single-precision word elements in vB by $2^{UIMM}$ . The products are converted to signed fixed-point integers using Round toward Zero mode. If the intermediate results are $> 2^{32}-1$ saturate to $2^{32}-1$ and if it is $< -2^{31}$ saturate to $-2^{31}$ . Place the signed integer results into the corresponding word elements of vD.

#### 4.2.2.5 Floating-Point Compare Instructions

This section describes floating-point unordered compare instructions.

All AltiVec floating-point compare instructions (**vcmpeqfp**, **vcmpgtfp**, **vcmpgefp**, and **vcmpbfp**) return FALSE if either operand is a NaN. Not equal-to, not greater-than, not greater-than-or-equal-to, and not-in-bounds NaNs compare to everything, including themselves.

Compares always return a Boolean mask (TRUE = 0xFFFF\_FFFF, FALSE = 0x0000\_0000) and never return a NaN. The **vcmpeqfp** instruction is recommended as the Isnan(vX) test. No explicit unordered compare instructions or traps are provided. However, the greater-than-or-equal-to predicate ( $\geq$ ) (**vcmpgefp**) is provided—in addition to the  $>$  and  $=$  predicates available for integer comparison—specifically to enable IEEE unordered comparison that would not be possible with just the  $>$  and  $=$  predicates. Table 4-10 lists the six common mathematical predicates and how they would be realized in AltiVec code.

Table 4-10. Common Mathematical Predicates

Case	Mathematical Predicate	AltiVec Realization	Relations			
			a>b	a<b	a=b	?
1	$a = b$	$a = b$	F	F	T	F
2	$a \neq b$ (?<>)	$\neg (a = b)$	T	T	F	T
3	$a > b$	$a > b$	T	F	F	F
4	$a < b$	$b > a$	F	T	F	F
5	$a \geq b$	$\neg (b > a)$	T	F	T	*T
6	$a \leq b$	$\neg (a > b)$	F	T	T	*T
5a	$a \geq b$	$a \geq b$	T	F	T	F
6a	$a \leq b$	$b \geq a$	F	T	T	F

\* Note: Cases 5 and 6 implemented with greater-than (**vcmpgfp** and **vnor**) would not yield the correct IEEE result when the relation is unordered.

Table Table 4-11 shows the remaining eight useful predicates and how they might be realized in AltiVec code.

Table 4-11. Other Useful Predicates

Case	Predicate	AltiVec Realization	Relations			
			a>b	a<b	a=b	?
7	$a ? b$	$\neg ((a=b) \vee (b>a) \vee (a>b))$	F	F	F	T
8	$a <> b$	$(a \geq b) \oplus (b \geq a)$	T	T	F	F
9	$a <=> b$	$(a \geq b) \vee (b \geq a)$	T	T	T	F
10	$a ?> b$	$\neg (b \geq a)$	T	F	F	T
11	$a ?>= b$	$\neg (b > a)$	T	F	T	T
12	$a ?< b$	$\neg (a \geq b)$	F	T	F	T
13	$a ?<= b$	$\neg (a > b)$	F	T	T	T
14	$a ?= b$	$\neg ((a > b) \vee (b > a))$	F	F	T	T

The vector floating-point compare instructions compare the elements in two vector registers word-by-word, interpreting the elements as single-precision numbers. With the exception of the Vector Compare Bounds Floating-Point (**vcmpbfp**) instruction they set the target vector register, and CR[6] if Rc = 1, in the same manner as do the vector integer compare instructions.

The Vector Compare Bounds Floating-Point (**vcmpbfp**) instruction sets the target vector register, and CR[6] if Rc = 1, to indicate whether the elements in **vA** are within the bounds specified by the corresponding element in **vB**, as explained in the instruction description. A

single-precision value  $x$  is said to be within the bounds specified by a single-precision value  $y$  if  $(-y \leq x \leq y)$ .

The floating-point compare instructions are summarized in Table 4-12.

**Table 4-12. Floating-Point Compare Instructions**

Name	Mnemonic	Syntax	Operation
Vector Compare Greater Than Floating-Point [Record]	vcmpgtfp[.]	vD,vA,vB	<p>Compare each of the four single-precision word elements in vA to the corresponding four single-precision word elements in vB</p> <p>For each element, if <math>vA &gt; vB</math> then set the corresponding element in vD to all 1's otherwise clear the element in vD to all 0's</p> <p>If the record bit is set (<math>Rc = 1</math>) in the vector compare instruction, then</p> <p>vD == 1, (all elements true) then CR6[0] is set</p> <p>vD == 0, (all elements false) then CR6[2] is set</p>
Vector Compare Equal to Floating-Point [Record]	vcmpeqfp[.]	vD,vA,vB	<p>Compare each of the 4 single-precision word elements in vA to the corresponding 4 single-precision word elements in vB.</p> <p>For each element, if <math>vA = vB</math> then set the corresponding element in vD to all 1's otherwise clear the element in vD to all 0's</p> <p>If the record bit is set (<math>Rc = 1</math>) in the vector compare instruction then</p> <p>vD == 1, (all elements true) then CR6[0] is set</p> <p>vD == 0, (all elements false) then CR6[2] is set</p>
Vector Compare Greater Than or Equal to Floating-Point [Record]	vcmpgefp[.]	vD,vA,vB	<p>Compare each of the 4 single-precision word elements in vA to the corresponding 4 single-precision word elements in vB.</p> <p>For each element, if <math>vA \geq vB</math> then set the corresponding element in vD to all 1's otherwise clear the element in vD to all 0's</p> <p>If the record bit is set (<math>Rc = 1</math>) in the vector compare instruction then</p> <p>vD == 1, (all elements true) then CR6[0] is set</p> <p>vD == 0, (all elements false) then CR6[2] is set</p>
Vector Compare Bounds Floating-Point [Record]	vcmpbfp[.]	vD,vA,vB	<p>Compare each of the 4 single-precision word elements in vA to the corresponding single-precision word elements in vB. A 2-bit value is formed that indicates whether the element in vA is within the bounds specified by the element in vB, as follows.</p> <p>Bit 0 of the two-bit value is cleared if the element in vA is <math>\leq</math> to the element in vB, and is set otherwise.</p> <p>Bit 1 of the two-bit value is cleared if the element in vA is <math>\geq</math> to the negation of the element in vB, and is set otherwise.</p> <p>The two-bit value is placed into the high-order two bits of the corresponding word element of vD and the remaining bits of the element are cleared to 0.</p> <p>If <math>Rc = 1</math>, CR6[2] is set when all four elements in vA are within the bounds specified by the corresponding element in vB</p>

### 4.2.2.6 Floating-Point Estimate Instructions

The floating-point estimate instructions are summarized in Table 4-13.

Table 4-13. Floating-Point Estimate Instructions

Name	Mnemonic	Syntax	Operation
Vector Reciprocal Estimate Floating-Point	vrefp	vD,vB	Place estimates of the reciprocal of each of the four word floating-point source elements in vB in the corresponding four word elements in vD.
Vector Reciprocal Square Root Estimate Floating-Point	vrsqrtefp	vD,vB	Place estimates of the reciprocal square-root of each of the four word source elements in vB in the corresponding four word elements in vD.
Vector Log2 Estimate Floating-Point	vlogefp	vD,vB	Place estimates of the base 2 logarithm of each of the four word source elements in vB in the corresponding four word elements in vD.
Vector 2 Raised to the Exponent Estimate Floating-Point	vexpteftp	vD,vB	Place estimates of 2 raised to the power of each of the four word source elements in vB in the corresponding four word elements in vD.

### 4.2.3 Load and Store Instructions

Only very basic load and store operations are provided in AltiVec ISA. This keeps the circuitry in the memory path fast so the latency of memory operations will be low. Instead, a powerful set of field manipulation instructions are provided to manipulate data into the desired alignment and arrangement after the data has been brought into the vector registers.

Load vector indexed (**lvx**, **lvxl**) and store vector indexed (**stvx**, **stvxl**) instructions transfer an aligned quad-word vector between memory and vector registers. Load vector element indexed (**lvebx**, **lvehx**, **lvewx**) and store vector element indexed instructions (**stvebx**, **stvehx**, **stvewx**) transfer byte, half-word, and word scalar elements between memory and vector registers.

All vector loads and vector stores use the index (**rA|0 + rB**) addressing mode to specify the target memory address. AltiVec ISA does not provide any update forms. An **lvebx**, **lvehx**, or **lvewx** instruction transfers a scalar data element from memory into the destination vector register, leaving other elements in the vector with boundedly-undefined values. A **stvebx**, **stvehx**, or **stvewx** instruction transfers a scalar data element from the source vector register to memory leaving other elements in the quad word unchanged. No data alignment occurs, that is, all scalar data elements are transferred directly on their natural memory byte-lanes to or from the corresponding element in the vector register. Quad word memory accesses made by **lvx**, **lvxl**, **stvx**, and **stvxl** instructions are not guaranteed to be atomic. Direct-store segments (T=1) are not supported by AltiVec ISA. Any vector load or store that attempts to access a direct-store segment will cause a DSI exception.

### 4.2.3.1 Alignment

All memory references must be size aligned. If a vector load or store address is not properly size aligned, the suitable number of least significant bits are ignored, and a size aligned transfer occurs instead. Data alignment must be performed by software after being brought into the registers. No assistance is provided for aligning individual scalar elements that are not aligned on their natural size boundary. However, assistance is provided for justifying non-size-aligned vectors. This is provided through the Load Vector for Shift Left (**lvsl**) and Load Vector for Shift Right (**lvslr**) instructions that compute the proper Vector Permute (**vperm**) control vector from the misaligned memory address. For details on how to use these instructions to align data see Section 3.1.6, “Quad-Word Data Alignment.”

The **lvx**, **lvxl**, **stvx**, and **stvxl** instructions can be used to move data, not just multimedia data, in PowerPC environments. Therefore, because vector loads and stores are size-aligned, care should be taken to align data on even quad-word boundaries for maximum performance.

### 4.2.3.2 Load and Store Address Generation

Vector load and store operations generate effective addresses using register indirect with index mode.

All AltiVec load and store instructions use register indirect with index addressing mode that cause the contents of two GPRs (specified as operands **rA** and **rB**) to be added in the generation of the effective address (EA). A zero in place of the **rA** operand causes a zero to be added to the value specified by **rB**. The option to specify **rA** or 0 is shown in the instruction descriptions as (**rA|0**). If the address becomes misaligned, for a half word, word, or quad word when combining addresses (**rA|0 + rB**), the effective address is ANDed with the appropriate zero values to boundary align the address and is summarized in Table 4-14.

**Table 4-14. Effective Address Alignment**

Operand	Effective Address Bit	Setting
Indexed half word	EA[63]	0b0
Indexed word	EA[62–63]	0b00
Indexed quad word	EA[60–63]	0b0000



Figure 4-1 shows how an effective address is generated when using register indirect with index addressing.

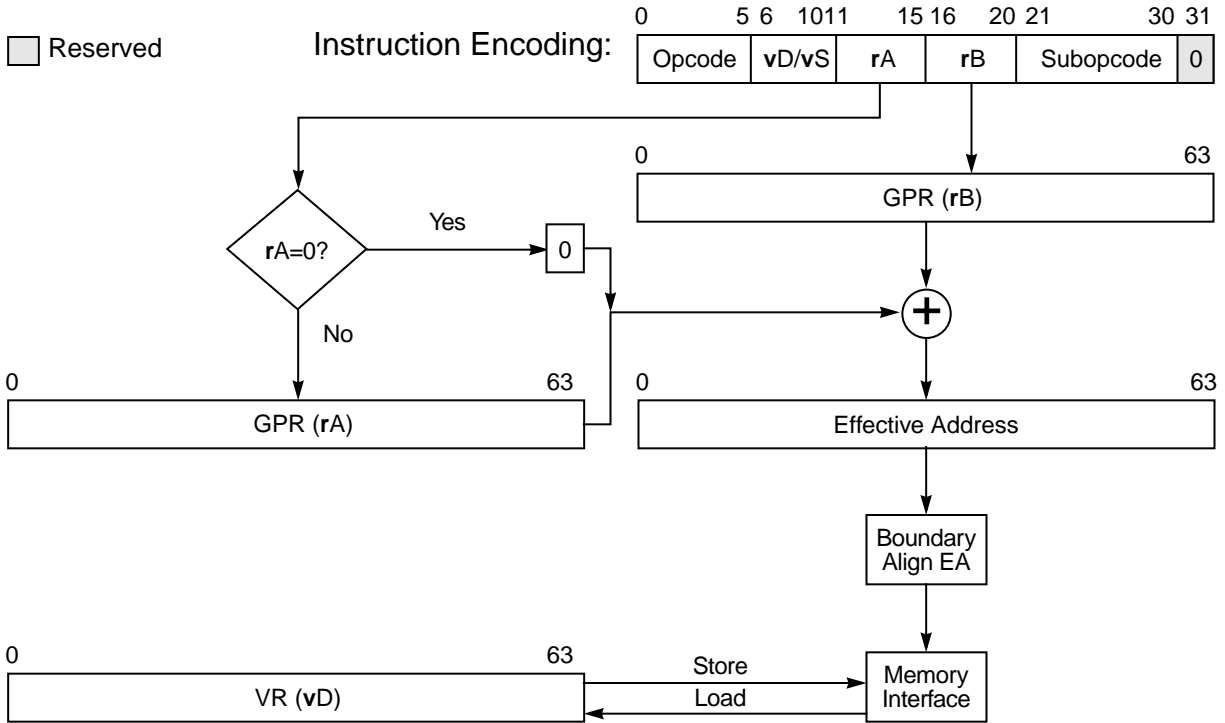


Figure 4-1. Register Indirect with Index Addressing for Loads/Stores

### 4.2.3.3 Vector Load Instructions

For vector load instructions, the byte, half word, or word addressed by the EA (effective address) is loaded into rD.

The default byte and bit ordering is big-endian as in the PowerPC architecture; see Section 3.1.2, “AltiVec Byte Ordering,” for information about little-endian byte ordering.

Table 4-15 summarizes the vector load instructions.

**Table 4-15. Integer Load Instructions**

Name	Mnemonic	Syntax	Operation
Load Vector Element Integer Indexed [b,h,w]	lvebx lvehx lvewx	vD,rA,rB	The EA is the sum (rA 0) + (rB). Load the byte, half word, or word in memory addressed by the EA into the low-order bits of vD. The remaining bits in vD are set to boundedly undefined values.  Because memory must stay aligned, the EA is set to default to alignment: For <b>b</b> , byte, integer length = 8 bits = 1 byte, For <b>h</b> , half word, integer length = 16 bits = 2 bytes, EA[62–63] is set to 0b0 For <b>w</b> , word, integer length = 32 bits = 4 bytes, EA[61-63] is set to 0b00
Load Vector Indexed	lvx	vD,rA,rB	The EA is the sum (rA 0) + (rB). Load the double word in memory addressed by the EA into vD.  Because memory needs to stay aligned, the EA is set to default to alignment:  For a quad word, integer length = 128 bits = 8 bytes, EA[60–63] is set to 0b0000  LRU = 0  If the processor is in little-endian mode, load the double word in memory addressed by EA into vD[64–127] and load the double word in memory addressed by EA+8 into vD[0–63].
Load Vector Indexed LRU	lvxl	vD,rA,rB	The EA is the sum (rA 0) + (rB). Load the double word in memory addressed by the EA into vD.  For the double word, integer length = 64 bits = 4 bytes, the EA[60–63] is set to 0b0000  LRU = 1, least recently used, hints that the quad word in the memory addressed by EA will probably not be needed again by the program in the near future.  If the processor is in little-endian mode, load the double word in memory addressed by EA into vD[64–127] and load the double word in memory addressed by EA+8 into vD[0–63].

The **lvsl** and **lvslr** instructions can be used to create the permute control vector to be used by a subsequent **vperm** instruction. Let X and Y be the contents of vA and vB specified by **vperm**. The control vector created by **lvsl** causes the **vperm** to select the high-order 16 bytes of the result of shifting the 32-byte value X || Y left by sh bytes (sh = the value in EA[60-63]). The control vector created by **lvslr** causes the **vperm** to select the low-order 16 bytes of the result of shifting X || Y right by sh bytes.

These instructions can also be used to rotate or shift the contents of a vector register left **lvsl** or right **lvslr** by sh bytes. The sh values for the **lvsl** instruction are shown in Table 4-17, and those for the **lvslr** instruction are shown in Table 4-18. For rotating, the vector register to be rotated should be specified as both the vA and the vB register for **vperm**. For shifting left, the vB register for **vperm** should be a register containing all zeros and vA should contain the value to be shifted, and vice versa for shifting right. For further examples on how to align the data see Section 3.1.6, “Quad-Word Data Alignment.” The default byte and bit

ordering is big-endian as in the PowerPC architecture; see Section 3.1.2.2, “Little-Endian Byte Ordering,” for information about little-endian byte ordering.

Table 4-16 summarizes the vector alignment instructions.

**Table 4-16. Vector Load Instructions Supporting Alignment**

Name	Mnemonic	Syntax	Operation
Load Vector for Shift Left	lvsl	vD,rA,rB	The EA is the sum (rA 0) + (rB). The EA[60–63] = sh, then based on Table 4-17, place the value in vD
Load Vector for Shift Right	lvsr	vD,rA,rB	The EA is the sum (rA 0) + (rB). The EA[60–63] = sh, then based on Table 4-18, place the value in vD

**Table 4-17. Shift Values for lvsl Instruction**

Shift (sh)	vD[0-127]
0x0	0x000102030405060708090A0B0C0D0E0F
0x1	0x0102030405060708090A0B0C0D0E0F10
0x2	0x02030405060708090A0B0C0D0E0F1011
0x3	0x0D0E0F101112131415161718191A1B1C
0x4	0x0405060708090A0B0C0D0E0F10111213
0x5	0x05060708090A0B0C0D0E0F1011121314
0x6	0x060708090A0B0C0D0E0F101112131415
0x7	0x0708090A0B0C0D0E0F10111213141516
0x8	0x08090A0B0C0D0E0F1011121314151617
0x9	0x090A0B0C0D0E0F101112131415161718
0xA	0x0A0B0C0D0E0F10111213141516171819
0xB	0x0B0C0D0E0F101112131415161718191A
0xC	0x0C0D0E0F101112131415161718191A1B
0xD	0x0D0E0F101112131415161718191A1B1C
0xE	0x0E0F101112131415161718191A1B1C1D
0xF	0x0F101112131415161718191A1B1C1D1E

**Table 4-18. Shift Values for lvsr Instruction**

Shift (sh)	vD[0-127]
0x0	0x101112131415161718191A1B1C1D1E1F
0x1	0x0F101112131415161718191A1B1C1D1E
0x2	0x0E0F101112131415161718191A1B1C1D
0x3	0x0D0E0F101112131415161718191A1B1C
0x4	0x0C0D0E0F101112131415161718191A1B

**Table 4-18. Shift Values for lvsr Instruction (continued)**

Shift (sh)	vD[0-127]
0x5	0x0B0C0D0E0F101112131415161718191A
0x6	0x0A0B0C0D0E0F10111213141516171819
0x7	0x090A0B0C0D0E0F101112131415161718
0x8	0x08090A0B0C0D0E0F1011121314151617
0x9	0x0708090A0B0C0D0E0F10111213141516
0xA	0x060708090A0B0C0D0E0F101112131415
0xB	0x05060708090A0B0C0D0E0F1011121314
0xC	0x0405060708090A0B0C0D0E0F10111213
0xD	0x030405060708090A0B0C0D0E0F101112
0xE	0x02030405060708090A0B0C0D0E0F1011
0xF	0x0102030405060708090A0B0C0D0E0F10

#### 4.2.3.4 Vector Store Instructions

For vector store instructions, the contents of vector register used as a source (vS) are stored into the byte, half word, word or quad word in memory addressed by the effective address (EA). Table 4-19 provides a summary of the vector store instructions.

**Table 4-19. Integer Store Instructions**

Name	Mnemonic	Syntax	Operation
Store Vector Element Indexed [b,h,w]	stvebx stvehx stvewx	vS,rA,rB	The EA is the sum (rA 0) + (rB). Store the contents of the low-order bits of vS into the integer in memory addressed by the EA.  Because memory needs to stay aligned, the EA is set to default to alignment:  For <b>b</b> , byte, integer length = 8 bits = 1 byte, For <b>h</b> , half word, integer length = 16 bits = 2 bytes, EA[62–63] is set to 0b0 For <b>w</b> , word, integer length = 32 bits = 4 bytes, EA[61–63] is set to 0b00

Table 4-19. Integer Store Instructions (continued)

Name	Mnemonic	Syntax	Operation
Store Vector Indexed	stvx	vS,rA,rB	<p>The EA is the sum (rA 0) + (rB). Store the contents of vS into the quad word in memory addressed by the EA.</p> <p>For <b>q</b>, quad word, integer length = 64 bits = 4 bytes, the EA[60–63] is set to 0b0000</p> <p>LRU = 0</p> <p>If the processor is in little-endian mode, store the contents of vS[64–127] into the double word in memory addressed by EA, and store the contents of vS[0–63] into the double word in memory addressed by EA+8.</p>
Store Vector Indexed LRU	stvxl	vD,rA,rB	<p>The EA is the sum (rA 0) + (rB). Store the contents of vS into the quad word in memory addressed by the EA.</p> <p>For <b>d</b>, double word, integer length=64 bits = 4 bytes, the EA[60–63] is set to 0b0000</p> <p>LRU = 1, least recently used, hints that the quad word in the memory addressed by EA will probably not be needed again by the program in the near future.</p> <p>If the processor is in little-endian mode, store the contents of vS[64–127] into the double word in memory addressed by EA, and store the contents of vS[0–63] into the double word in memory addressed by EA+8.</p>

## 4.2.4 Control Flow

AltiVec instructions can be freely intermixed with existing PowerPC instructions to form a complete program. AltiVec instructions do provide a vector compare and select mechanism to implement conditional execution as a mechanism to control data flow in AltiVec programs. And AltiVec vector compare instructions can update the condition register thus providing the communication from AltiVec execution units to PowerPC branch instructions necessary to modify program flow based on vector data.

## 4.2.5 Vector Permutation and Formatting Instructions

Vector pack, unpack, merge, splat, permute, and select can be used to accelerate various vector math and vector formatting. Details of the various instructions follow.

### 4.2.5.1 Vector Pack Instructions

Half-word vector pack instructions (**vpkuhum**, **vpkuhus**, **vpkshus**, **vpkshss**) truncate the sixteen half words from two concatenated source operands producing a single result of sixteen bytes (quad word) using either modulo( $2^8$ ), 8-bit signed-saturation, or 8-bit unsigned-saturation to perform the truncation. Similarly, word vector pack instructions (**vpkuwum**, **vpkuwus**, **vpkswus**, and **vpksws**) truncate the eight words from two concatenated source operands producing a single result of eight half words using

modulo( $2^{16}$ ), 16-bit signed-saturation, or 16-bit unsigned-saturation to perform the truncation.

One special form of Vector Pack Pixel (**vpkpx**) instruction packs eight 32-bit (8/8/8/8) pixels from two concatenated source operands into a single result of eight 16-bit 1/5/5/5  $\alpha$ RGB pixels. The least significant bit of the first 8-bit element becomes the 1-bit  $\alpha$  field, and each of the three 8-bit R, G, and B fields are reduced to 5 bits by ignoring the 3 lsbs.

Table 4-20 describes the vector pack instructions.

**Table 4-20. Vector Pack Instructions**

Name	Mnemonic	Syntax	Operation
Vector Pack Unsigned Integer [h,w] Unsigned Modulo	vpkuhum vpkuwum	vD, vA, vB	Concatenate the low-order unsigned integers of vA and the low-order unsigned integers of vB and place into vD using unsigned modulo arithmetic. vA is placed in the lower order double word of vD and vB is placed into the higher order double word of vD.  For h, half word, integer length = 16 bits = 2 bytes, eight unsigned integers, in other words the 8 low-order bytes of the half words from vA and vB  For w, word, integer length = 32 bits = 4 bytes, four unsigned integers, in other words the 4 low-order half words of the words from vA and vB
Vector Pack Unsigned Integer [h,w] Unsigned Saturate	vpkuhus vpkuwus	vD, vA, vB	Concatenate the low-order unsigned integers of vA and the low-order unsigned integers of vB and place into vD using unsigned saturate clamping mode. vA is placed in the lower order double word of vD and vB is placed into the higher order double word of vD.  For h, half word, integer length = 16 bits = 2 bytes, eight unsigned integers, in other words the 8 low-order bytes of the half words from vA and vB  For w, word, integer length = 32 bits = 4 bytes, four unsigned integers, in other words the 4 low-order words of the half words from vA and vB
Vector Pack Signed Integer [h,w] Unsigned Saturate	vpkshus vpkswus	vD, vA, vB	Concatenate the low-order signed integers of vA and the low-order signed integers of vB and place into vD using unsigned saturate clamping mode. vA is placed in the lower order double word of vD and vB is placed into the higher order double word of vD.  For h, half word, integer length = 16 bits = 2 bytes, eight signed integers, in other words the 8 low-order bytes of the half word from vA and vB  For w, word, integer length = 32 bits = 4 bytes, four signed integers, in other words the 4 low-order half words of the words from vA and vB
Vector Pack Signed Integer [h,w] Signed Saturate	vpkshss vpkswss	vD, vA, vB	Concatenate the low-order signed integers of vA and the low-order signed integers of vB are concatenated and place into vD using signed saturate clamping mode. vA is placed in the lower order double word of vD and vB is placed into the higher order double word of vD.  For h, half word, integer length = 16 bits = 2 bytes, eight signed integers, in other words the 8 low-order bytes of the half word from vA and vB  For w, word, integer length = 32 bits = 4 bytes, four signed integers, in other words the 4 low-order half words of the words from vA and vB

Table 4-20. Vector Pack Instructions (continued)

Name	Mnemonic	Syntax	Operation
Vector Pack Pixel	vpx	vD, vA, vB	<p>Each word element in vA and vB is packed to 16 bits and the half word is placed into vD. Each word from vA and vB is packed to 16 bits in the following order:</p> <p>[bit 7 of the first byte (bit 7 of the word)]</p> <p>[bits 0–4 of the second byte (bits 8–12 of the word)]</p> <p>[bits 0–4 of the third byte (bits 16–20 of the word)]</p> <p>[bits 0–4 of the fourth byte (bits 24–28 of the word)]</p> <p>vA half words are placed in the lower order double word of vD and vB half words are placed into the higher order double word of vD.</p> <p>For h, half word, integer length = 16 bits = 2 bytes, eight signed integers, in other words the 8 low-order bytes of the half word from vA and vB</p> <p>For w, word, integer length = 32 bits = 4 bytes, four signed integers, in other words the 4 low-order half words of the words from vA and vB</p>

### 4.2.5.2 Vector Unpack Instructions

Byte vector unpack instructions unpack the 8 low bytes (or 8 high bytes) of one source operand into 8 half words using sign extension to fill the MSBs. Half word vector unpack instructions unpack the 4 low half words (or 4 high half words) of one source operand into 4 words using sign extension to fill the MSBs.

A special purpose form of vector unpack is provided, the Vector Unpack Low Pixel (**vupklpx**) and the Vector Unpack High Pixel (**vupkhp**) instructions for 1/5/5/5  $\alpha$ RGB pixels. The 1/5/5/5 pixel vector unpack, unpacks the four low 1/5/5/5 pixels (or four 1/5/5/5 high pixels) into four 32-bit (8/8/8/8) pixels. The 1-bit  $\alpha$  element in each pixel is sign extended to 8 bits, and the 5-bit R, G, and B elements are each zero extended to 8 bits.

Table 4-21 describes the unpack instructions.

Table 4-21. Vector Unpack Instructions

Name	Mnemonic	Syntax	Operation
Vector Unpack High Signed Integer [b,h]	vupkhsb vupkshs	vD, vB	Each signed integer element in the high order double word of vB is sign extended to fill the MSBs in a signed integer and then is placed into vD.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, eight signed bytes from the high order double word of vB are unpacked and sign extended to 8 half words into vD.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, eight signed half words from the high order double word of vB are unpacked and sign extended to 4 words into vD
Vector Unpack High Pixel	vupkhp	vD, vB	Each half-word element in the high order double word of vB is unpacked to produce a 32-bit word that is then placed in the same order into vD.  A half-word element is unpacked to 32 bits by concatenating, in order, the results of the following operations.  sign-extend bit 0 of the half word to 8 bits zero-extend bits 1–5 of the half word to 8 bits zero-extend bits 6–10 of the half word to 8 bits zero-extend bits 11–15 of the half word to 8 bits
Vector Unpack Low Signed Integer [b,h]	vupklsb vupklsh	vD, vB	Each signed integer element in the low-order double word of vB is sign extended to fill the MSBs in a signed integer and then is placed into vD.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, eight signed bytes from the low-order double word of vB are unpacked and sign extended to 8 half words into vD.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, eight signed half words from the low-order double word of vB are unpacked and sign extended into 4 words in vD
Vector Unpack Low Pixel	vupklp	vD, vB	Each half-word element in the low-order double word of vB is unpacked to produce a 32-bit word that is then placed in the same order into vD.  A half-word element is unpacked to 32 bits by concatenating, in order, the results of the following operations.  sign-extend bit 0 of the half word to 8 bits zero-extend bits 1–5 of the half word to 8 bits zero-extend bits 6–10 of the half word to 8 bits zero-extend bits 11–15 of the half word to 8 bits

### 4.2.5.3 Vector Merge Instructions

Byte vector merge instructions interleave the 8 low bytes (or 8 high bytes) from two source operands producing a result of 16 bytes. Similarly, half-word vector merge instructions interleave the 4 low half words (or 4 high half words) of two source operands producing a result of 8 half words, and word vector merge instructions interleave the 2 low words (or 2 high words) from two source operands producing a result of 4 words. The vector merge instruction has many uses, notable among them is a way to efficiently transpose SIMD vectors. Table 4-22 describes the merge instructions.



**Table 4-22. Vector Merge Instructions**

Name	Mnemonic	Syntax	Operation
Vector Merge High Integer [b,h,w]	vmrghb vmrghh vmrghw	vD, vA, vB	<p>Each integer element in the high order double word of vA is placed into the low-order integer element in vD. Each integer element in the high order double word of vB is placed into the high order integer element in vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, 8 bytes from the high order double word of vA are placed into the low-order byte of each half word in vD and 8 bytes from the high order double word of vB are placed into the high order byte of each half word in vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 4 half words from the high order double word of vA are placed into the low-order half word of each word in vD and 4 half words from the high order double word of vB are placed into the high order half word of each word in vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 2 words from the high order double word of vA are placed into the low-order word of each double word in vD and 2 words from the high order double word of vB are placed into the high order word of each double word in vD.</p>
Vector Merge Low Integer [b,h,w]	vmrglb vmrglh vmrglw	vD, vA, vB	<p>Each integer element in the low-order double word of vA is placed into the low-order integer element in vD. Each integer element in the low-order double word of vB is placed into the high order integer element in vD.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, 8 bytes from the low-order double word of vA are placed into the low-order byte of each half word in vD and 8 bytes from the low-order double word of vB are placed into the high order byte of each half word in vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 4 half words from the low-order double word of vA are placed into the low-order half word of each word in vD and 4 half words from the low-order double word of vB are placed into the high order half word of each word in vD.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 2 words from the low-order double word of vA are placed into the low-order word of each double word in vD and 2 words from the low-order double word of vB are placed into the high order word of each double word in vD.</p>

#### 4.2.5.4 Vector Splat Instructions

When a program needs to perform arithmetic vector, the vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a Vector Register by a constant). Vector splat instructions can be used to move data where it is required. For example to multiply all elements of a vector register by a constant, the vector splat instructions can be used to splat the scalar into the vector register. Likewise, when storing a scalar into an arbitrary memory location, it must be splatted into a vector register, and that register must be specified as the source of the store. This will guarantee that the data appears in all possible positions of that scalar size for the store. Table 4-23 describes the vector splat instructions.

Table 4-23. Vector Splat Instructions

Name	Mnemonic	Syntax	Operation
Vector Splat Integer [b,h,w]	vspltb vsplth vspltw	vD, vB, UIMM	Replicate the contents of element UIMM in vB and place into each element in vD. For <b>b</b> , byte, integer length = 8 bits = 1 byte, each element is a byte. For <b>h</b> , half word, integer length = 16 bits = 2 bytes, each element is a half word. For <b>w</b> , word, integer length = 32 bits = 4 bytes, 2 words each element is a word.
Vector Splat Immediate Signed Integer [b,h,w]	vspltsb vspltsih vspltsiw	vD, SIMM	Sign-extend the value of the SIMM field to the length of the element and replicate that value and place into each element in vD. For <b>b</b> , byte, integer length = 8 bits = 1 byte, each element is a byte. For <b>h</b> , half word, integer length = 16 bits = 2 bytes, each element is a half word. For <b>w</b> , word, integer length = 32 bits = 4 bytes, 2 words each element is a word.

#### 4.2.5.5 Vector Permute Instruction

Permute instructions allow any byte in any two source vector registers to be directed to any byte in the destination vector. The fields in a third source operand specify from which field in the source operands the corresponding destination field will be taken. The Vector Permute (**vperm**) instruction is a very powerful one that provides many useful functions. For example, it provides a good way to perform table-lookups and data alignment operations. An example of how to use the command in aligning data see Section 3.1.6, “Quad-Word Data Alignment.” Table 4-24 describes the vector permute instruction.

Table 4-24. Vector Permute Instruction

Name	Mnemonic	Syntax	Operation
Vector Permute	vperm	vD, vA,vB,vC	vC specifies which bytes from vA and vB are to be copied and placed into the byte elements in vD.

#### 4.2.5.6 Vector Select Instruction

Data flow in the vector unit can be controlled without branching by using a vector compare and the vector select (**vsel**) instructions. In this use, the compare result vector is used directly as a mask operand to vector select instructions. The **vsel** instruction selects one field from one or the other of two source operands under control of its mask operand. Use of the TRUE/FALSE compare result vector with select in this manner produces a two instruction equivalent of conditional execution on a per-field basis. Table 4-25 describes the **vsel** instruction.

Table 4-25. Vector Select Instruction

Name	Mnemonic	Syntax	Operation
Vector Select	vsel	vD,vA,vB,vC	For each bit, compare the value in vC to the value 0b0 and if it equals 0b0 then load vD with vA's corresponding bit value otherwise compare the value in vC to the value 0b1 and if it equals 0b1 then load vD with vB's corresponding bit value.

### 4.2.5.7 Vector Shift Instructions

The vector shift instructions shift the contents of a vector register or of a pair of vector registers left or right by a specified number of bytes (**vslo**, **vsro**, **vsldoi**) or bits (**vsl**, **vsr**). Depending on the instruction, this shift count is specified either by low-order bits of a vector register or by an immediate field in the instruction. In the former case the low-order 7 bits of the shift count register give the shift count in bits ( $0 \leq \text{count} \leq 127$ ). Of these 7 bits, the high-order 4 bits give the number of complete bytes by which to shift and are used by **vslo** and **vsro**; the low-order 3 bits give the number of remaining bits by which to shift and are used by **vsl** and **vsr**.

There are two methods of specifying an inter-element shift or rotate of two source vector registers, extracting 16 bytes as the result vector. There is also a method for shifting a single source vector register left or right by any number of bits.

Table 4-26 describes the various vector shift instructions.

Table 4-26. Vector Shift Instructions

Name	Mnemonic	Syntax	Operation
Vector Shift Left	vsl	vD,vA,vB	Shift vA left by the 3 lsbs of vB, and place the result into vD If vB value in invalid, the default result is boundedly undefined
Vector Shift Right	vsr	vD,vA,vB	Shift vA right by the 3 lsbs of vB, and place the result into vD If vB value in invalid, the default result is boundedly undefined
Vector Shift Left Double by Octet Immediate	vsldoi	vD,vA,vB,SH	Shift vB left by the 3 lsbs of SH value and then OR with vA, place the result is into vD If vB value in invalid, the default result is 0
Vector Shift Left by Octet	vslo	vD,vA,vB	Shift vA left by the 3 lsbs of vB, and place the result into vD If vB value in invalid, the default result is 0b000
Vector Shift Right by Octet	vsro	vD,vA,vB	Shift vA right by the 3 lsbs of vB, and place the result into vD If vB value in invalid, the default result is 0b000

#### 4.2.5.7.1 Immediate Interelement Shifts/Rotates

The Vector Shift Left Double by Octet Immediate (**vsldoi**) instruction provides the basic mechanism that can be used to provide inter-element shifts and/or rotates. This instruction is like a **vperm**, except that the shift count is specified as a literal in the instruction rather

than as a control vector in another vector register, as is required by **vperm**. The result vector consists of the left-most 16 bytes of the rotated 32-byte concatenation of **vA:vB**, where shift (SH) is the rotate count. Table 4-27 below enumerates how various shift functions can be achieved using the **vsidoi** instruction.

**Table 4-27. Coding Various Shifts and Rotates with the vsidoi Instruction**

To Get This:		Code This:			
Operation	sh	Instruction	Immediate	vA	vB
Rotate left double	0–15	<b>vsidoi</b>	0–15	MSV	LSV
Rotate left double	16–31	<b>vsidoi</b>	mod16(SH)	LSV	MSV
Rotate right double	0–15	<b>vsidoi</b>	16–sh	MSV	LSV
Rotate right double	16–31	<b>vsidoi</b>	16–mod16(SH)	LSV	MSV
Shift left single, zero fill	0–15	<b>vsidoi</b>	0–15	MSV	0x0
Shift right single, zero fill	0–15	<b>vsidoi</b>	16–SH	0x0	MSV
Rotate left single	0–15	<b>vsidoi</b>	0–15	MSV	=vA
Rotate right single	0–15	<b>vsidoi</b>	16–SH	MSV	=vA

#### 4.2.5.7.2 Computed Interelement Shifts/Rotates

The Load Vector for Shift Left (**lvsl**) instruction and Load Vector for Shift Right (**lvsr**) instruction are supplied to assist in shifting and/or rotating vector registers by an amount determined at run time. The input specifications have the same form as the vector load and store instructions, that is, it uses register indirect with index addressing mode(**rA|0 +rB**). This is because one of their primary purposes is to compute the permute control vector necessary for post-load and pre-store shifting necessary for dealing with misaligned vectors.

This **lvsl** instruction can be used to align a big-endian misaligned vector after loading the (aligned) vectors that contain its pieces. The **lvsl** instruction can be used to misalign a vector register for use in a read-modify-write sequence that will store an misaligned little-endian vector.

The **lvsr** instruction can be used to align a little-endian misaligned vector after loading the (aligned) vectors that contain its pieces. The **lvsl** instruction can be used to misalign a vector register for use in a read-modify-write sequence that will store an misaligned big-endian vector.

For an example on how the **lvsl** instruction is used to align a vector in big-endian mode see Section 3.1.6.1, “Accessing a Misaligned Quad Word in Big-Endian Mode.” For an example on how **lvsr** is used to align a vector in little-endian mode see Section 3.1.6.2, “Accessing a Misaligned Quad Word in Little-Endian Mode.”

### 4.2.5.7.3 Variable Interelement Shifts

A vector register may be shifted left or right by a number of bits specified in a vector register. This operation is supported with four instructions, two for right shift and two for left shift.

The Vector Shift Left by Octet (**vslo**) and Vector Shift Right by Octet (**vsro**) instructions shift a vector register from 0 to 15 bytes as specified in bits 121–124 of another vector register. The Vector Shift Left (**vsl**) and Vector Shift Right (**vsr**) instructions shift a vector register from 0 to 7 bits as specified in another vector register (the shift count must be specified in the three lsbs of each byte in the vector and must be identical in all bytes or the result is boundedly undefined). In all of these instructions, zeros are shifted into vacated element and bit positions.

Used sequentially with the same shift-count vector register, these instructions will shift a vector register left or right from 0 to 127 bits as specified in bits 121–127 of the shift-count vector register. For example:

```
vslo      VZ, VX, VY
vspltb   VY, VY, 15
vsl      VZ, VZ, VY
```

will shift **vX** by the number of bits specified in **vY** and place the results in **vZ**.

With these instructions a full double-register shift can be performed in seven instructions. The following code will shift **vW||vX** left by the number of bits specified in **vY** placing the result in **vZ**:

```
vslo      t1, VW, VY      ; shift the most significant. register left
vspltb   VY, VY, 15
vsl      t1, t1, VY
vsububm  VY, V0, VY      ; adjust count for right shift (V0=0)
vsro     t2, VX, VY      ; right shift least sign. register
vsr      t2, t2, VY
vor      VZ, t1, t2      ; merge to get the final result
```

## 4.2.6 Processor Control Instructions—UISA

Processor control instructions are used to read from and write to the PowerPC condition register (CR), machine state register (MSR), and special-purpose registers (SPRs). See Chapter 4, “Addressing Mode and Instruction Set Summary,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*, for information about the instructions used for reading from and writing to the MSR and SPRs.

### 4.2.6.1 AltiVec Status and Control Register Instructions

Table 4-28 summarizes the instructions for reading from or writing to the Vector Status and Control Register (VSCR). For more information on VSCR see section in Section 2.3.2, “Vector Status and Control Register (VSCR).”

**Table 4-28. Move to/from Condition Register Instructions**

Name	Mnemonic	Syntax	Operation
Move to Vector Status and Control Register	<b>mtvscr</b>	CRM,rS	Place the contents of vB into VSCR.
Move from Vector Status and Control Register	<b>mfvscr</b>	vB	Place the contents of VSCR into vB.

### 4.2.7 Recommended Simplified Mnemonics

To simplify assembly language programs, a set of simplified mnemonics is provided for some of the most frequently used operations (such as no-op, load immediate, load address, move register, and complement register). Assemblers could provide the simplified mnemonics listed below. Programs written to be portable across the various assemblers for PowerPC architecture should not assume the existence of mnemonics not described in this document.

Simplified mnemonics are provided for the Data Stream Touch (**dst**) and Data Stream Touch for Store (**dstst**) instructions so that they can be coded with the transient indicator as part of the mnemonic rather than as a numeric operand. Similarly, simplified mnemonics are provided for the Data Stream Stop (**dss**) instruction so that it can be coded with the all streams indicator is part of the mnemonic. These are shown as examples with the instructions in Table 4-29.

**Table 4-29. Simplified Mnemonics for Data Stream Touch (dst)**

Operation	Simplified Mnemonic	Equivalent to
Data Stream Touch (non-transient)	<b>dst</b> rA, rB, STRM	<b>dst</b> rA, rB, STRM,0
Data Stream Touch Transient	<b>dstt</b> rA, rB, STRM	<b>dst</b> rA, rB, STRM,1
Data Stream Touch for Store (non-transient)	<b>dstst</b> rA, rB, STRM	<b>dstst</b> rA, rB, STRM,0
Data Stream Touch for Transient	<b>dststt</b> rA, rB, STRM	<b>dststt</b> rA, rB, STRM,1
Data Stream Stop (one stream)	<b>dss</b> STRM	<b>dss</b> STRM,0
Data Stream Stop All	<b>dssall</b>	<b>dss</b> 0,1

## 4.3 AltiVec VEA Instructions

- U** PowerPC virtual environment architecture (VEA) describes the semantics of the memory model that can be assumed by software processes, and includes descriptions of the cache model, cache-control instructions, address aliasing, and other related issues.
- V**
- ©**

Implementations that conform to the VEA also adhere to the UISA, but may not necessarily adhere to the OEA. For further details see Chapter 4, “Addressing Mode and Instruction Set Summary,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

This section describes the additional AltiVec instructions defined for the VEA.

### 4.3.1 Memory Control Instructions—VEA

Memory control instructions include the following types:

- Cache management instructions (user-level and supervisor-level)
- Segment register manipulation instructions
- Segment lookaside buffer management instructions
- Translation lookaside buffer (TLB) management instructions

This section describes the user-level cache management instructions defined by the VEA. See Chapter 4, “Addressing Mode and Instruction Set Summary,” in *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture* for more information about supervisor-level cache, segment register manipulation, and TLB management instructions.

### 4.3.2 User-Level Cache Instructions—VEA

The instructions summarized in this section provide user-level programs the ability to manage on-chip caches if they are implemented. See Chapter 5, “Cache Model and Memory Coherency,” in *The Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture* for more information about cache topics.

Bandwidth between the processor and memory is managed explicitly by the programmer through the use of cache management instructions. These instructions give software a way to communicate to the cache hardware how it should prefetch and prioritize writeback of data. The principal instruction for this purpose is a software directed cache prefetch instruction called Data Stream Touch (**dst**). Other related instructions are provided for complete control of the software directed cache prefetch mechanism.

Table 4-30 summarizes the directed prefetch cache instructions defined by the VEA. Note that these instructions are accessible to user-level programs. See Section 5.2.1, “Software-Directed Prefetch for further details on the prefetch cache instructions.

Table 4-30. User-Level Cache Instructions

Name	Mnemonic	Syntax	Operation
Data Stream Touch	<b>dst</b>	rA,rB,STRM,T	<p>This instruction associates the data stream specified by the contents of rA and rB with the stream ID specified by STRM.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (rA), where rA ≠ 0                      unit size: (rB)[3–7] if (rB)[3–7] ≠ 0; otherwise 32                      count: (rB)[8–15] if (rB)[8–15] ≠ 0; otherwise 256                      stride: (rB)[16–31] if (rB)[16–31] ≠ 0; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1).</p> <p>If rA=0, the instruction form is invalid.</p> <p>See Section 5.2.1.1, “Data Stream Touch (dst),” for further details on the <b>dst</b> instruction.</p>
Data Stream Touch	<b>dstt</b>	rA,rB,STRM,T	<p>This instruction associates the data stream specified by the contents of registers rA and rB with the stream ID specified by STRM.</p> <p>This instruction is a hint that performance will probably be improved if the cache blocks containing the specified data stream are not fetched into the data cache, because the program will probably not load from the stream. That is, the data stream will be relatively transient in nature. That is, it will have poor locality and is likely to be referenced a very few times or over a very short period of time. The memory subsystem can use this persistent/transient knowledge to manage the data as is most appropriate for the specific design of the cache/memory hierarchy of the processor on which the program is executing. An implementation is free to ignore <b>dstt</b>, in that case it should simply be executed as a <b>dst</b>. However, software should always attempt to use the correct form of <b>dst</b> or <b>dstt</b> regardless of whether the intended processor implements <b>dstt</b>. In this way the program will automatically benefit when run on processors that support <b>dstt</b>.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (rA), where rA ≠ 0                      unit size: (rB)[3–7] if (rB)[3–7] ≠ 0; otherwise 32                      count: (rB)[8–15] if (rB)[8–15] ≠ 0; otherwise 256                      stride: (rB)[16–31] if (rB)[16–31] ≠ 0; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be accessed into fairly frequently in the near future (T=0) or to be transient (T=1).</p> <p>If rA=0, the instruction form is invalid.</p> <p>See Section 5.2.1.2, “Transient Streams,” for further details on the <b>dstt</b> instruction.</p>



**Table 4-30. User-Level Cache Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Data Stream Touch for Store (non-transient)	<b>dstst</b>	rA,rB,STRM,T	<p>This instruction associates the data stream specified by the contents of registers rA and rB with the stream ID specified by STRM.</p> <p>This instruction is a hint that performance will probably be improved if the cache blocks containing the specified data stream are fetched into the data cache, because the program will probably soon access into the stream, and that prefetching from any data stream that was previously associated with the specified stream ID is no longer needed. The hint is ignored for blocks that are caching inhibited.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (rA), where rA ≠ 0  unit size: (rB)[3–7] if (rB)[3-7] ≠ 0; otherwise 32  count: (rB)[8–15] if (rB)[8–15] ≠ 0; otherwise 256  stride: (rB)[16–31] if (rB)[16–31] ≠ 0; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1).</p> <p>If rA=0, the instruction form is invalid.</p> <p>See Section 5.2.1.3, “Storing to Streams (dstst),” for further details on the <b>dstst</b> instruction.</p>
Data Stream Touch for Store	<b>dststt</b>	rA,rB,STRM,T	<p>This instruction associates the data stream specified by the contents of rA and rB with the stream ID specified by STRM.</p> <p>This instruction is a hint that performance will probably not be improved if the cache blocks containing the specified data stream are fetched into the data cache, because the program will probably not access the stream. That is, the data stream will be relatively transient in nature. That is, it will have poor locality and is likely to be referenced a very few times or over a very short period of time. The memory subsystem can use this persistent/transient knowledge to manage the data as is most appropriate for the specific design of the cache/memory hierarchy of the processor on which the program is executing.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (rA), where rA ≠ 0  unit size: (rB)[3–7] if (rB)[3-7] ≠ 0; otherwise 32  count: (rB)[8–15] if (rB)[8–15] ≠ 0; otherwise 256  stride: (rB)[16–31] if (rB)[16–31] ≠ 0; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1).</p> <p>If rA=0, the instruction form is invalid.</p> <p>See Section 5.2.1.3, “Storing to Streams (dstst),” for further details on the <b>dststt</b> instruction.</p>

**Table 4-30. User-Level Cache Instructions (continued)**

Name	Mnemonic	Syntax	Operation
Data Stream Stop	<b>dss</b>	STRM,A	<p>If A = 0 and a data stream associated with the stream ID specified by STRM exists, this instruction terminates prefetching of that data stream.</p> <p>If A = 1, this instruction terminates prefetching of all existing data streams. (The STRM field is ignored.)</p> <p>In addition, executing a <b>dss</b> instruction ensures that all memory accesses associated with data stream prefetching caused by preceding <b>dst</b> and <b>dstst</b> instructions that specified the same stream ID as that specified by the <b>dss</b> instruction (A = 0), or by all preceding <b>dst</b> and <b>dstst</b> instructions (A = 1), will be in group G1 with respect to the memory barrier created by a subsequent <b>sync</b> instruction.</p> <p><b>dss</b> serves as both a basic and an extended mnemonic. The assembler will recognize a <b>dss</b> mnemonic with two operands as the basic form, and a <b>dss</b> mnemonic with one operand as the extended form.</p> <p>Execution of a <b>dss</b> instruction causes address translation for the specified data stream(s) to cease. Prefetch requests for which the effective address has already been translated may complete and may place the corresponding data into the data cache</p> <p>See Section 5.2.1.4, "Stopping Streams," for further details on the <b>dss</b> instruction.</p>
Data Stream Stop All	<b>dssall</b>		<p>Terminates prefetching of all existing data streams. All active streams may be stopped.</p> <p>If the optional data stream prefetch facility is implemented, <b>dssall</b> (extended mnemonic for <b>dss</b>), to terminate any data stream prefetching requested by the interrupted program, in order to avoid prefetching data in the wrong context, consuming memory bandwidth fetching data that are not likely to be needed by the other program, and interfering with data cache use by the other program. The <b>dssall</b> must be followed by a <b>sync</b>, and additional software synchronization may be required.</p> <p>See Section 5.2.1.4, "Stopping Streams," for further details on the <b>dssall</b> instruction.</p>

# Chapter 5

## Cache, Exceptions, and Memory Management

This chapter summarizes details of AltiVec™ technology that pertain to cache and memory management models. Note that AltiVec technology defines most of its instructions at the user level (UISA). Because most AltiVec instructions are computational, there is little effect on the VEA and OEA portions of the PowerPC architecture definition.

Because the AltiVec instruction set architecture (ISA) uses 128-bit operands, additional instructions are provided to optimize cache and memory bus use.

### 5.1 PowerPC Shared Memory ▼

To fully understand the data stream prefetch instructions for AltiVec, one needs a knowledge of PowerPC architecture for shared memory. The PowerPC architecture supports the sharing of memory between programs, between different instances of the same program, and between processors and other mechanisms. It also supports access to memory by one or more programs using different effective addresses. All these cases are considered memory sharing. Memory is shared in blocks that are an integral number of pages.

When the same memory has different effective addresses, the addresses are called aliases. Each application can be granted separate access privileges to aliased pages. For more details on how the PowerPC architecture supports the sharing of memory see Chapter 5, “Cache Model and Memory Coherency” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 5.2 AltiVec Memory Bandwidth Management ■

The AltiVec ISA provides a way for software to speculatively load larger blocks of data from memory. That is, bandwidth otherwise idle can be used to permit software to take advantage of locality and reduces the number of system memory accesses.

## 5.2.1 Software-Directed Prefetch

Bandwidth between the processor and memory is managed explicitly by the programmer using cache management instructions. These instructions let software indicate to the cache hardware how to prefetch and prioritize data writeback. The principle instruction for this purpose is a software-directed cache prefetch instruction, Data Stream Touch (**dst**), described in the following section.

### 5.2.1.1 Data Stream Touch (**dst**)


The data stream prefetch facility permits a program to indicate that a sequence of units of memory is likely to be accessed soon by memory access instructions. Such a sequence is called a data stream or, when the context is clear, simply a stream. A data stream is defined by the following:

- EA—The effective address of the first unit in the sequence
- Unit size—The number of quad words in each unit;  $0 < \text{unit size} \leq 32$
- Count—The number of units in the sequence;  $0 < \text{count} \leq 256$
- Stride—The number of bytes between the effective address of one unit in the sequence and the effective address of the next unit in the sequence (that is, the effective address of the *n*th unit in the sequence is  $\text{EA} + (n - 1) \times \text{stride}$ );  $(-32768 \leq \text{stride} < 0$  or  $0 < \text{stride} \leq 32768)$

The units need not be aligned on a particular memory boundary. The stride may be negative.

The **dst** instruction specifies a starting address, a block size (1–32 vectors), a number of blocks to prefetch (1–256 blocks), and a signed stride in bytes (-32,768 to +32,768 bytes). The 2-bit tag, specified as an immediate field in the opcode, identifies one of four possible touch streams. The starting address of the stream is specified in **rA** (if **rA** = 0, the instruction form is invalid). BlockSize, BlockCount, and BlockStride are specified in **rB**. Do not confuse the term ‘cache block’: the term ‘block’ always indicates a PowerPC cache block.

The format of the **rB** register is shown in Figure 5-1.



**Figure 5-1. Format of rB in dst Instruction**

There is no zero-length block size, block count, or block stride. A BlockSize of 0 indicates 32 vectors, a BlockCount of 0 indicates 256 blocks, and a BlockStride of 0 indicates +32,768 bytes. Otherwise, these fields correspond to the numerical value of the size, count, and stride. Do not specify strides smaller than 1 block (16 bytes).

The programmer specifies block size in terms of vectors (16 bytes), regardless of the cache-block size. Hardware automatically optimizes the number of cache blocks it fetches to bring a block into the cache. The number of cache blocks fetched into the cache for each block is the fewest natural cache blocks needed to fetch the entire block, including the effects of block misalignment to cache blocks, as shown in the following:

$$\text{CacheBlocksFetched} = \text{ceiling} \left( \frac{\text{BlockSize} + \text{mod}(\text{BlockAddr}, \text{CacheBlockSize})}{\text{CacheBlockSize}} \right)$$

The address of each block in a stream is a function of the stream's starting address, the block stride, and the block being fetched. The starting address may be any 32-bit byte address. Each block's address is computed as a full 32-bit byte address from the following:

$$\text{BlockAddr}_n = (\text{rA}) + n (\text{rB})_{16-31} \quad \text{where } n = \{0 \dots (\text{BlockCount} - 1)\}$$

and if  $((\text{rB})_{16-31} = 0)$  then  $((\text{rB})_{16-31}) \Leftarrow 32768$

The address of the first cache block fetched in each block is that block's address aligned to the next lower natural cache-block boundary by ignoring  $\log_2(\text{CacheBlockSize})$  least significant bits (lsbs) (for example, for 32-byte cache-blocks, the five lsbs are ignored). Cache blocks are then fetched sequentially forward until the entire block of vectors is brought into the cache. An example of a six-block data stream is shown in Figure 5-2

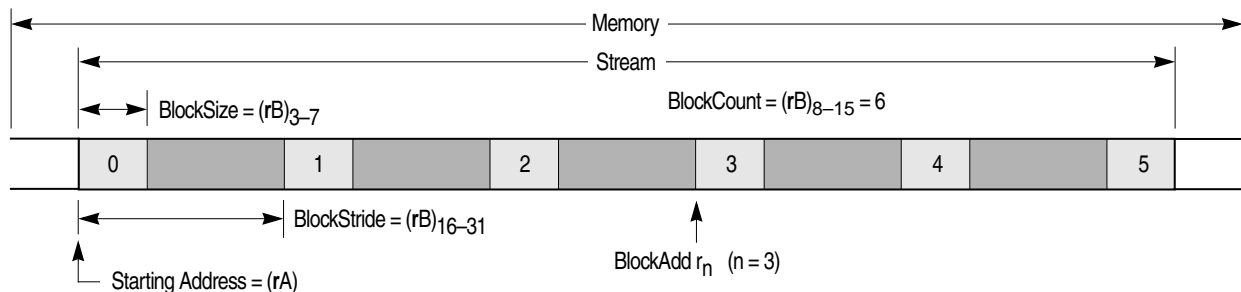


Figure 5-2. Data Stream Touch

Executing a **dst** instruction notifies the cache/memory subsystem that the program will soon need specified data. If bandwidth is available, the hardware starts loading the specified stream into the cache. To the extent that hardware can acquire the data, when the loads requiring the data finally execute, the target data will be in the cache. Executing a second **dst** to the tag of a stream in progress aborts the existing stream (at hardware's earliest convenience) and establishes a new stream with the same stream tag ID.

The **dst** instruction is a hint to hardware and has no architecturally visible effects (in the PowerPC UISA sense). The hardware is free to ignore it, to start the prefetch when it can, to abort the stream at any time, or to prioritize other memory operations ahead of it. If a stream is aborted, the program still functions properly, but subsequent loads experience the full latency of a cache miss.

The **dst** instruction does not introduce implementation problems like those of load/store multiple/string instructions. Because **dst** does not affect the architectural state, it does not cause interlock problems associated with load/store multiple/string instructions. Also, **dst** does take exceptions and requires no complex recovery mechanism.

Touch instructions should be considered strong hints. Using them in highly speculative situations could waste considerable bandwidth. Implementations that do not implement the stream mechanism treat stream instructions (**dst**, **dstt**, **dsts**, **dstst**, **dss**, and **dssall**) as no-ops. If the stream mechanism is implemented, all four streams must be provided.

### 5.2.1.2 Transient Streams

The memory subsystem considers **dst** an indication that its stream data is likely to have some reasonable degree of locality and be referenced several times or over some reasonably long period. This is called persistence. The Data Stream Touch Transient instruction (**dstt**) indicates to the memory system that its stream data is transient, that is, it has poor locality and is likely to be used very few times or only for a very short time. A memory subsystem can use this knowledge to manage data for the processor's cache/memory design. An implementation may ignore the distinction between transience and persistence; in that case, **dstt** acts like **dst**. However, portable software should always use the correct form of **dst** or **dstt** regardless of whether the intended processor makes that distinction.

### 5.2.1.3 Storing to Streams (dstst)

A **dst** instruction brings a cache block into the cache subsystem in a state most efficient for subsequent reading of data from it (load). The companion instruction, Data Stream Touch for Store (**dstst**), brings the cache block into the cache subsystem in a state most efficient for subsequent writing to it (store). For example, in a MESI cache subsystem, a **dst** might bring a cache block in shared (S) state, whereas a **dstst** would bring the cache block in exclusive (E) state to avoid a subsequent demand-driven bus transaction to take ownership of the cache block so the store can proceed.

The **dstst** streams are the same physical streams as **dst** streams, that is, **dstst** stream tags are aliases of **dst** tags. If not implemented, **dstst** defaults to **dst**. If **dst** is not implemented, it is a no-op. The **dststt** instruction is a transient version of **dstst**.

Data stream prefetching of memory locations is not supported when bit 57 of the segment table entry or bit 0 of the segment register (SR) is set. If a **dst** or **dstst** instruction specifies a data stream containing these memory locations, results are undefined.

### 5.2.1.4 Stopping Streams

The **dst** instructions have a counterpart called Data Stream Stop (**dss**). A program can stop any given stream prefetch by executing **dss** with that stream's tag. This is useful when a program speculatively starts a stream prefetch but later determines that the instruction stream went the wrong way. The **dss** instruction can stop the stream so that no more bandwidth is wasted. All active streams may be stopped by using **dssall**. This is useful when the operating system needs to stop all active streams (process switch), but does not know how many streams are in progress.

Because **dssall** does not specify the number of implemented streams, it should always be used instead of a sequence of **dss** instructions to stop all streams.

Neither **dss** nor **dssall** is execution synchronizing; the time between when a **dss** is issued and the stream stops is not specified. Therefore, when software must ensure that the stream is physically stopped before continuing (for example, before changing virtual memory mapping), a special sequence of synchronizing instructions is required. The sequence can differ for different situations, but the following sequence works in all contexts:

```
dssall          ; stop all streams
sync           ; insert a barrier in memory pipe
lwz           Rn,... ; stick one more operation in memory pipe
cmpd          Rn,Rn ;
bne-         *-4   ; make sure load data is back
isync         ; wait for all previous instructions to
               ; complete to ensure
               ; memory pipe is clear and nothing is
               ; pending in the old context
```

Data stream prefetching for a given stream is terminated by executing the appropriate **dss** instruction. The termination can be synchronized by executing a **sync** instruction after the **dss** instruction if the memory barrier created by **sync** orders all address translation effects of the subsequent context-altering instructions. Otherwise, data dependencies are also required. For example, the following instruction sequence terminates all data stream prefetching before altering the contents of an segment register (SR):

```
dssall          ; stop all data stream prefetching
sync           ; order dssall before load
lwz           Ry,sr_y(Rx); load new SR value
mtsr          y,Ry ; alter rY
```

The **mtsr** instruction cannot be executed until the **lwz** loads the SR value into **rY**. The memory access caused by the **lwz** cannot be performed until the **dssall** instruction takes effect (that is, until address translation stops for all data streams and all memory accesses associated with data stream prefetches for which the effective address was translated before the translation stops are performed).

### 5.2.1.5 Exception Behavior of Prefetch Streams

In general, exceptions do not cancel streams. Streams are sensitive to whether the processor is in user or supervisor mode (determined by MSR[PR]) and whether data address translation is used (determined by MSR[DR]). This allows prefetch streams to behave predictably when an exception occurs.

Streams are suspended in real addressing mode (MSR[DR] = 0) and remain suspended until translation is turned back on (MSR[DR] is set). A **dst** instruction issued while MSR[DR] = 0 produces boundedly undefined results.

A stream is suspended whenever the MSR[PR] is different from what it was when the **dst** that established it was issued. For example, if a **dst** is issued in user mode (MSR[PR] = 1), the resulting stream is suspended when the processor enters supervisor mode (MSR[PR] = 0) and remains suspended until the processor returns to user mode. Conversely, if the **dst** were issued in supervisor mode, it is suspended if the machine enters user mode.

Because exceptions do not cancel streams automatically, the operating system must stop streams explicitly when warranted, for example, when switching processes or changing virtual memory context. Care must be taken if data stream prefetching is used in supervisor-level state (MSR[PR] = 0).

After an exception is taken, the supervisor-level program that next changes MSR[DR] from 0 to 1 causes data-stream prefetching to resume for any data streams for which the corresponding **dst** or **dstst** instruction was executed in supervisor mode; such streams are called supervisor-level data streams. This program is unlikely to be the one that executed the corresponding **dst** or **dstst** instruction and is unlikely to use the same address translation context as that in which the **dst** or **dstst** was executed. Suspension and resumption of data stream prefetching work more naturally for user level data streams, because the next application program to be dispatched after an exception occurs is likely to be the most recently interrupted program. An exception handler that changes the context in which data addresses are translated may need to terminate data-stream prefetching for supervisor-level data streams and to synchronize the termination before changing MSR[DR] to 1.

Although terminating all data stream prefetching in this case would satisfy the requirements of the architecture, doing so would adversely affect the performance of applications that use data-stream prefetching. Thus, it may be better for the operating system to record stream IDs associated with any supervisor-level data streams and to terminate prefetching for those streams only.

Cache effects of supervisor-level data-stream prefetching can also adversely affect performance of applications that use data stream prefetching, as supervisor-level use of the associated stream ID can take over an application's data stream.

Data stream instructions cannot cause exceptions directly. Therefore, any event that would cause an exception on a normal load or store, such as a page fault or protection violation, is instead aborted and ignored.



Suspension or termination of data stream prefetching for a given data stream need not cancel prefetch requests for that data stream for which the effective address has been translated and need not cause data returned by such requests to be discarded. However, to improve software's ability to pace data stream prefetching with data consumption, it may be better to limit the number of these pending requests that can exist simultaneously.

### 5.2.1.6 Synchronization Behavior of Streams

Streams are not affected (stopped or suspended) by execution of any PowerPC synchronization instructions (**sync**, **isync**, or **eieio**). This permits these instructions to be used for synchronizing multiple processors without disturbing background prefetch streams. Prefetch streams have no architecturally observable effects and are not affected by synchronization instructions. Synchronizing the termination of data stream prefetching is needed only by the operating system

### 5.2.1.7 Address Translation for Streams

Like **dcbt** and **dcbtst** instructions, **dst**, **dstst**, **dstt**, and **dststt** are treated as loads with respect to address translation, memory protection, and reference and change recording.

Unlike **dcbt** and **dcbtst** instructions, stream instructions that cause a TLB miss cause a page table search and the page descriptor to be loaded into the TLB. Conceptually, address translation and protection checking is performed on every cache-block access in the stream and proceeds normally across page boundaries and TLB misses, terminating only on page faults or protection violations that cause a DSI exception.

Stream instructions operate like normal PowerPC cache instructions (such as **dcbt**) with respect to guarded memory; they are not subject to normal restrictions against prefetching in guarded space because they are program-directed. However, speculative **dst** instructions can not start a prefetch stream to guarded space.

If the effective address of a cache block within a data stream cannot be translated, or if loading from the block would violate memory protection, the processor will terminate prefetching of that stream. (Continuing to prefetch subsequent cache blocks within the stream might cause prefetching to get too far ahead of consumption of prefetched data.) If the effective address can be translated, a TLB miss can cause such termination, even on implementations for which TLBs are reloaded in software.

### 5.2.1.8 Stream Usage Notes

A given data stream exists if a **dst** or **dstst** instruction has been executed that specifies the stream and prefetching of the stream has neither completed, terminated, or been supplanted. Prefetching of the stream has completed, when all the memory locations within the stream that will ever be prefetched as a result of executing the **dst** or **dstst** instruction have been prefetched (for example, locations for which the effective address cannot be translated will

never be prefetched). Prefetching of the stream is terminated by executing the appropriate **dss** instruction; it is supplanted by executing another **dst** or **dstst** instruction that specifies the stream ID associated with the given stream. Because there are four stream IDs, as many as four data streams may exist simultaneously.

The maximum block count of **dst** is small because of its preferred usage. It is not intended for a single **dst** instruction to prefetch an entire data stream. Instead, **dst** instructions should be issued periodically, for example on each loop iteration, for the following reasons:

- Short, frequent **dst** instructions better synchronize the stream with the consumption of data.
- With prefetch closely synchronized just ahead of consumption, another activity is less likely to inadvertently evict prefetched data from the cache before it is needed.
- The prefetch stream is restarted automatically after an exception (that could have caused the stream to be terminated by the operating system) with no additional complex hardware mechanisms needed to restart the prefetch stream.

Issuing new **dst** instructions to stream tag IDs in progress terminates old streams—**dst** instructions cannot be queued.

For example, when multiple **dst** instructions are used to prefetch a large stream, it would be poor strategy to issue a second **dst** whose stream begins at the specified end of the first stream before it was certain that the first stream had completed. This could terminate the first stream prematurely, leaving much of the stream unprefetched.

Paradoxically, it would also be unwise to wait for the first stream to complete before issuing the second **dst**. Detecting completion of the first stream is not possible, so the program would have to introduce a pessimistic waiting period before restarting the stream and then incur the full start-up latency of the second stream.

The correct strategy is to issue the second **dst** well before the anticipated completion of the first stream and begin it at an address overlapping the first stream by an amount sufficient to cover any portion of the first stream that could not yet have been prefetched. Issuing the second **dst** too early is not a concern because blocks prefetched by the first stream hit in the cache and need not be refetched. Thus, even if issued prematurely and overlapped excessively, the second **dst** rapidly advances to the point of prefetching new blocks. This strategy allows a smooth transition from the first stream to the second without significant breaks in the prefetch stream.

For the greatest performance benefit from data-stream prefetching, use the **dst** and **dstst** (and **dss**) instructions so that the prefetched data is used soon after it is available in the data cache. Pacing data stream prefetching with consumption increases the likelihood that prefetched data is not displaced from the cache before it is used, and reduces the likelihood that prefetched data displaces other data needed by the program.

Specifying each logical data stream as a sequence of shorter data streams helps achieve the desired pacing, even in the presence of exceptions, and address translation failures. The components of a given logical data stream should have the following attributes:

- The same stream ID should be associated with each component.
- The components should partially overlap (that is, the first part of a component should consist of the same memory locations as the last part of the preceding component).
- The memory locations that do not overlap with the next component should be large enough that a substantial portion of the component is prefetched. That is, prefetch enough memory locations for the current component before it is taken over by the prefetching being done for the next component.

### 5.2.1.9 Stream Implementation Assumptions

Some processors can treat **dst** instructions as no-ops. However, if a processor implements **dst**, a minimum level of functionality is provided to create as consistent a programming model across different machines as possible. A program can assume the following functionality in a **dst** instruction:

- Implements all four tagged streams
- Implements each tagged stream as a separate, independent stream with arbitration for memory access performed on a round-robin basis.
- Searches the table for each stream access that misses in the TLB.
- Does not abort streams on page boundary crossings
- Does not abort streams on exceptions (except DSI exceptions caused by the stream).
- Does not abort streams, or delay execution pending completion of streams, on PowerPC synchronization instructions **sync**, **isync**, or **eieio**.
- Does not abort streams on TLB misses that occur on loads or stores issued concurrently with running streams. However, a DSI exception from one of those loads or stores may cause streams to abort.

### 5.2.2 Prioritizing Cache Block Replacement

Load Vector Indexed LRU (**lvxl**) and Store Vector Indexed LRU (**stvxl**) instructions provide explicit control over cache block replacement by letting the programmer indicate whether an access is likely to be the last reference made to the cache block containing this load or store. The cache hardware can then prioritize replacement of this cache block over others with older but more useful data.

Data accessed by a normal load or store is likely to be needed more than once. Marking this data as most-recently used (MRU) indicates that it should be a low-priority candidate for

replacement. However, some data, such as that used in DSP multimedia algorithms, is rarely reused and should be marked as the highest priority candidate for replacement.

Normal accesses mark data MRU. Data unlikely to be reused can be marked LRU. For example, on replacing a cache block marked LRU by one of these instructions, a processor may improve cache performance by evicting the cache block without storing it in intermediate levels of the cache hierarchy (except to maintain cache consistency).

### 5.2.3 Partially Executed AltiVec Instructions

The OEA permits certain instructions to be partially executed when an alignment or DSI exception occurs. In the same way that the target register may be altered when floating-point load instructions cause a DSI exception, if the AltiVec facility is implemented, the target register (**vD**) may be altered when **lvx** or **lvxl** is executed and the TLB entry is invalidated before the access completes.

Exceptions cause data stream prefetching to be suspended for all existing data streams. Prefetching for a given data stream resumes when control is returned to the interrupted program, if the stream still exists (for example, the operating system did not terminate prefetching for the stream).

## 5.3 DSI Exception—Data Address Breakpoint

A data address breakpoint register (DABR) match causes a DSI exception in implementations that support the data breakpoint feature. When a DABR match occurs on a non-AltiVec processor that support the PowerPC architecture, the DAR is set to any effective address between and including the word (for a byte, half word, or word access) specified by the effective address computed by the instruction and the effective address of the last byte in the word or double word in which the match occurred. In processors that support the AltiVec technology, this would include a quad-word access from an **lvx**, **lvxl**, **stvx**, or **stvxl** instruction to a segment or BAT area.


## 5.4 AltiVec Unavailable Exception (0x00F20)

The AltiVec facility includes an additional instruction-caused, precise exception to those defined by the OEA and discussed in Chapter 6, “Exceptions,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*. An AltiVec unavailable exception occurs when no higher priority exception exists (see Table 5-2), an attempt is made to execute an AltiVec instruction, and **MSR[VEC] = 0**.

Register settings for AltiVec unavailable exceptions are described in Table 5-1 and shown in Figure 5-3.

**Table 5-1. AltiVec Unavailable Exception—Register Settings**

Register	Setting Description																																								
SRR0	Set to the effective address of the instruction that caused the exception																																								
SRR1	<p><b>32-Bit</b>                      0Loaded with equivalent bits from the MSR                      1–4Cleared                      5–9Loaded with equivalent bits from the MSR                      10–15Cleared                      16–31 Loaded with equivalent bits from the MSR                      Note that depending on the implementation, additional MSR bits may be copied to SRR1.</p>																																								
MSR	<table border="0"> <tr> <td>SF</td><td>1</td><td>EE</td><td>0</td><td>SE</td><td>0</td><td>DR</td><td>0</td> </tr> <tr> <td>ISF</td><td>—</td><td>PR</td><td>0</td><td>BE</td><td>0</td><td>RI</td><td>0</td> </tr> <tr> <td>VEC</td><td>0</td><td>FP</td><td>0</td><td>FE1</td><td>0</td><td>LE</td><td>Set to value of ILE</td> </tr> <tr> <td>POW</td><td>0</td><td>ME</td><td>—</td><td>IP</td><td>—</td><td></td><td></td> </tr> <tr> <td>ILE</td><td>—</td><td>FE0</td><td>0</td><td>IR</td><td>0</td><td></td><td></td> </tr> </table>	SF	1	EE	0	SE	0	DR	0	ISF	—	PR	0	BE	0	RI	0	VEC	0	FP	0	FE1	0	LE	Set to value of ILE	POW	0	ME	—	IP	—			ILE	—	FE0	0	IR	0		
SF	1	EE	0	SE	0	DR	0																																		
ISF	—	PR	0	BE	0	RI	0																																		
VEC	0	FP	0	FE1	0	LE	Set to value of ILE																																		
POW	0	ME	—	IP	—																																				
ILE	—	FE0	0	IR	0																																				



**Figure 5-3. SRR1 Bit Settings after an AltiVec Unavailable Exception**

When an AltiVec unavailable exception is taken, instruction execution resumes as offset 0x00F20 from the base address determined by MSR[IP].

The **dst** and **dstst** instructions are supported if MSR[DR] = 1. If either instruction is executed when MSR[DR] = 0 (real addressing mode), results are boundedly undefined.

Conditions that cause this exception are prioritized among instruction-caused (synchronous), precise exceptions as shown in Table 5-2, taken from the section “Exception Priorities,” in Chapter 6, “Exceptions,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

**Table 5-2. Exception Priorities (Synchronous/Precise Exceptions)**

Priority	Exception
3 <sup>1</sup>	<p>Instruction dependent—When an instruction causes an exception, the exception mechanism waits for any instructions prior to the excepting instruction in the instruction stream to complete. Any exceptions caused by these instructions are handled first. It then generates the appropriate exception if no higher priority exception exists when the exception is to be generated.</p> <p>Note that a single instruction can cause multiple exceptions. When this occurs, those exceptions are ordered in priority as indicated in the following:</p> <p>A. Integer loads and stores</p> <ul style="list-style-type: none"> <li>a. Alignment</li> <li>b. DSI</li> <li>c. Trace (if implemented)</li> </ul> <p>B. Floating-point loads and stores</p> <ul style="list-style-type: none"> <li>a. Floating-point unavailable</li> <li>b. Alignment</li> <li>c. DSI</li> <li>d. Trace (if implemented)</li> </ul> <p>C. Other floating-point instructions</p> <ul style="list-style-type: none"> <li>a. Floating-point unavailable</li> <li>b. Program—Precise-mode floating-point enabled exception</li> <li>c. Floating-point assist (if implemented)</li> <li>d. Trace (if implemented)</li> </ul> <p>D. AltiVec loads and stores (if AltiVec facility implemented)</p> <ul style="list-style-type: none"> <li>a. AltiVec unavailable</li> <li>b. DSI</li> <li>c. Trace (if implemented)</li> </ul> <p>E. Other AltiVec Instructions (if AltiVec facility implemented)</p> <ul style="list-style-type: none"> <li>a. AltiVec unavailable</li> <li>b. Trace (if implemented)</li> </ul> <p>F. The <b>rfi</b> and <b>mtmsr</b></p> <ul style="list-style-type: none"> <li>a. Program—Supervisor level Instruction</li> <li>b. Program—Precise-mode floating-point enabled exception</li> <li>c. Trace (if implemented), for <b>mtmsr</b> only</li> </ul> <p>If precise-mode IEEE floating-point enabled exceptions are enabled and FPSCR[FEX] is set, a program exception occurs no later than the next synchronizing event.</p> <p>G. Other instructions</p> <ul style="list-style-type: none"> <li>a. These exceptions are mutually exclusive and have the same priority: <ul style="list-style-type: none"> <li>— Program: Trap</li> <li>— System call (<b>sc</b>)</li> <li>— Program: Supervisor level instruction</li> <li>— Program: Illegal Instruction</li> </ul> </li> <li>b. Trace (if implemented)</li> </ul> <p>F. ISI exception</p> <p>The ISI exception has the lowest priority in this category. It is only recognized when all instructions prior to the instruction causing this exception appear to have completed and that instruction is to be executed. The priority of this exception is specified for completeness and to ensure that it is not given more favorable treatment. An implementation can treat this exception as though it had a lower priority.</p>

<sup>1</sup> The exceptions are third in priority after system reset and machine check exceptions

# Chapter 6

## AltiVec Instructions

This chapter lists the AltiVec instruction set in alphabetical order by mnemonic. Note that each entry includes the instruction format and a graphical representation of the instruction. All the instructions are 32 bit and a description of the instruction fields and pseudocode conventions are also provided. For more information on the AltiVec instruction set, refer to Chapter 4 “Addressing Modes and Instruction Set Summary.” For more information on the PowerPC instruction set, refer to Chapter 8, “Instruction Set,” in the *Programming Environments Manual for 32-Bit Implementations of the PowerPC Architecture*.

### 6.1 Instruction Formats

AltiVec instructions are four bytes (32 bits) long and are word-aligned. AltiVec instruction set architecture (ISA) has four operands, three source vectors, and one result vector. Bits 0–5 always specify the primary opcode for AltiVec instructions. AltiVec ALU-type instructions specify the primary opcode point 4 (0b00\_01\_00). AltiVec load, store, and stream prefetch instructions use secondary opcode in primary opcode 31 (0b01\_11\_11).

Within a vector register, a byte, half-word, or word element are referred to as follows:

- Byte elements, each byte = 8 bits; in the pseudocode,  $n = 8$  with a total of 16 elements
- Half-word elements, each byte = 16 bits; in the pseudocode,  $n = 16$  with a total of 8 elements
- Word elements, each byte = 32 bits; in the pseudocode,  $n = 32$  with a total of 4 elements

Refer to Figure 1-3 for an example of how elements are placed in a vector register.

#### 6.1.1 Instruction Fields

Table 6-1 describes the instruction fields used in the various instruction formats.

**Table 6-1. Instruction Syntax Conventions**

Field	Description
OPCD (0–5)	Primary opcode field
rA, A (11–15)	Specifies a GPR to be used as a source or destination
rB, B (16–20)	Specifies a GPR to be used as a source
Rc (31)	Record bit 0 Does not update the condition register (CR). 1 For the optional AltiVec facility, set CR field 6 to control program flow as described in Section 2.4.1, “PowerPC Condition Register”
vA (11–15)	Specifies a vector register to be used as a source
vB (16–20)	Specifies a vector register to be used as a source
vC (21–25)	Specifies a vector register to be used as a source
vD (6–10)	Specifies a vector register to be used as a destination
vS (6–10)	Specifies a vector register to be used as a source
SHB (22–25)	Specifies a shift amount in bytes.
SIMM (11–15)	This immediate field is used to specify a (5-bit) signed integer.
UIMM (11–15)	This immediate field is used to specify a 4-, 8-, 12-, or 16-bit unsigned integer.

## 6.1.2 Notation and Conventions

The operation of some instructions is described by a semiformal language (pseudocode). See Table 6-2 for a list of additional pseudocode notation and conventions used throughout this section.

**Table 6-2. Notation and Conventions**

Notation/Convention	Meaning
←	Assignment
¬	NOT logical operator
do i=X to Y by Z	Do the following starting at X and iterating to Y by Z
+ <sub>int</sub>	2's complement integer add
¬ <sub>int</sub>	2's complement integer subtract
+ <sub>ui</sub>	Unsigned integer add
¬ <sub>ui</sub>	Unsigned integer subtract
* <sub>ui</sub>	Unsigned integer multiply
+ <sub>si</sub>	Signed integer add
¬ <sub>si</sub>	Signed integer subtract
* <sub>si</sub>	Signed integer multiply



Table 6-2. Notation and Conventions (continued)

Notation/Convention	Meaning
* <sub>sui</sub>	Signed integer (first operand) multiplied by unsigned integer (second operand) producing signed result
/	Integer divide
+ <sub>fp</sub>	Single-precision floating-point add
- <sub>fp</sub>	Single-precision floating-point subtract
* <sub>fp</sub>	Single-precision floating-point multiply
÷ <sub>fp</sub>	Single-precision floating-point divide
√ <sub>fp</sub>	Single-precision floating-point square root
< <sub>ui</sub> , ≤ <sub>ui</sub> , > <sub>ui</sub> , ≥ <sub>ui</sub>	Unsigned integer comparison relations
< <sub>si</sub> , ≤ <sub>si</sub> , > <sub>si</sub> , ≥ <sub>si</sub>	Signed integer comparison relations
< <sub>fp</sub> , ≤ <sub>fp</sub> , > <sub>fp</sub> , ≥ <sub>fp</sub>	Single precision floating point comparison relations
≠	Not equal
= <sub>int</sub>	Integer equal to
= <sub>ui</sub>	Unsigned integer equal to
= <sub>si</sub>	Signed integer equal to
= <sub>fp</sub>	Floating-point equal to
X >> <sub>ui</sub> Y	Shift X right by Y bits extending Xs vacated bits with zeros
X >> <sub>si</sub> Y	Shift X right by Y bits extending Xs vacated bits with the sign bit of X
X << <sub>ui</sub> Y	Shift X left by Y bits inserting Xs vacated bits with zeros
	Used to describe the concatenation of two values (that is, 010    111 is the same as 010111)
&	AND logical operator
	OR logical operator
⊕, ≡	Exclusive-OR, Equivalence logical operators (for example, (a ≡ b) = (a ⊕ ¬ b))
0bnnnn	A number expressed in binary format.
0xnnnn	A number expressed in hexadecimal format.
?	Unordered comparison relation
X <sub>0</sub>	X zeros
X <sub>1</sub>	X ones
X <sub>Y</sub>	X copies of Y
X <sub>Y</sub>	bit Y of X
X <sub>Y:Z</sub>	bits Y through Z, inclusive, of X
LENGTH(x)	Length of x, in bits. If x is the word “elemen,” LENGTH(x) is the length, in bits, of the element implied by the instruction mnemonic.

**Table 6-2. Notation and Conventions (continued)**

Notation/Convention	Meaning
ROTL(x,y)	Result of rotating x left by y bits
UtoUImod(X,Y)	Chop unsigned integer X- to Y-bit unsigned integer
UtoUlsat(X,Y)	Result of converting the unsigned-integer x to a y-bit unsigned-integer with unsigned-integer saturation
SlttoUlsat(X,Y)	Result of converting the signed-integer x to a y-bit unsigned-integer with unsigned-integer saturation
SlttoSImod(X,Y)	Chop integer X- to Y-bit integer
SlttoSlsat(X,Y)	Result of converting the signed-integer x to a y-bit signed-integer with signed-integer saturation
RndToNearFP32	The single-precision floating-point number that is nearest in value to the infinitely-precise floating-point intermediate result x (in case of a tie, the even single-precision floating-point value is used).
RndToFPInt32Near	The value x if x is a single-precision floating-point integer; otherwise the single-precision floating-point integer that is nearest in value to x (in case of a tie, the even single-precision floating-point integer is used).
RndToFPInt32Trunc	The value x if x is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than x if x>0, or the smallest single-precision floating-point integer that is greater than x if x<0
RndToFPInt32Ceil	The value x if x is a single-precision floating-point integer; otherwise the smallest single-precision floating-point integer that is greater than x
RndToFPInt32Floor	The value x if x is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than x
CnvtFP32ToUI32Sat(x)	Result of converting the single-precision floating-point value x to a 32-bit unsigned-integer with unsigned-integer saturation
CnvtFP32ToSI32Sat(x)	Result of converting the single-precision floating-point value x to a 32-bit signed-integer with signed-integer saturation
CnvtUI32ToFP32(x)	Result of converting the 32-bit unsigned-integer x to floating-point single format
CnvtSI32ToFP32(x)	Result of converting the 32-bit signed-integer x to floating-point single format
MEM(X,Y)	Value at memory location X of size Y bytes
SwapDouble	Swap the doublewords in a quadword vector
ZeroExtend(X,Y)	Zero-extend X on the left with zeros to produce Y-bit value
SignExtend(X,Y)	Sign-extend X on the left with sign bits (that is, with copies of bit 0 of x) to produce Y-bit value
RotateLeft(X,Y)	Rotate X left by Y bits
mod(X,Y)	Remainder of X/Y
UImaximum(X,Y)	Maximum of 2 unsigned integer values, X and Y
SImaximum(X,Y)	Maximum of 2 signed integer values, X and Y
FPmaximum(X,Y)	Maximum of 2 floating-point values, X and Y
UIminimum(X,Y)	Minimum of 2 unsigned integer values, X and Y

Table 6-2. Notation and Conventions (continued)

Notation/Convention	Meaning
SIminimum(X,Y)	Minimum of 2 unsigned integer values, X and Y
FPminimum(X,Y)	Minimum of 2 floating-point values, X and Y
FPReciprocalEstimate12(X)	12-bit-accurate floating-point estimate of 1/X
FPReciprocalSQREstimate12(X)	12-bit-accurate floating-point estimate of 1/(sqrt(X))
FPLog <sub>2</sub> Estimate3(X)	3-bit-accurate floating-point estimate of log <sub>2</sub> (X)
FPPower2Estimate3(X)	3-bit-accurate floating-point estimate of 2**X
CarryOut(X + Y)	Carry out of the sum of X and Y
ROTL[64](x, y)	Result of rotating the 64-bit value x left y positions
ROTL[32](x, y)	Result of rotating the 32-bit value x    x left y positions, where x is 32 bits long
0bnnnn	A number expressed in binary format.
0xnxxx	A number expressed in hexadecimal format.
(n)x	The replication of x, n times (that is, x concatenated to itself n – 1 times). (n)0 and (n)1 are special cases. A description of the special cases follows: <ul style="list-style-type: none"> <li>• (n)0 means a field of n bits with each bit equal to 0. Thus (5)0 is equivalent to 0b00000.</li> <li>• (n)1 means a field of n bits with each bit equal to 1. Thus (5)1 is equivalent to 0b11111.</li> </ul>
(rA)0	The contents of rA if the rA field has the value 1–31, or the value 0 if the rA field is 0.
(rX)	The contents of rX
x[n]	n is a bit or field within x, where x is a register
x <sup>n</sup>	x is raised to the n <sup>th</sup> power
ABS(x)	Absolute value of x
CEIL(x)	Least integer ≥ x
Characterization	Reference to the setting of status bits in a standard way that is explained in the text.
CIA	Current instruction address. The 32-bit address of the instruction being described by a sequence of pseudocode. Used by relative branches to set the next instruction address (NIA) and by branch instructions with LK = 1 to set the link register. Does not correspond to any architected register.
Clear	Clear the leftmost or rightmost n bits of a register to 0. This operation is used for rotate and shift instructions.
Clear left and shift left	Clear the leftmost b bits of a register, then shift the register left by n bits. This operation can be used to scale a known non-negative array index by the width of an element. These operations are used for rotate and shift instructions.
Cleared	Bits = 0.

**Table 6-2. Notation and Conventions (continued)**

Notation/Convention	Meaning
Do	Do loop. <ul style="list-style-type: none"> <li>• Indenting shows range.</li> <li>• “To” and/or “by” clauses specify incrementing an iteration variable.</li> <li>• “While” clauses give termination conditions.</li> </ul>
DOUBLE(x)	Result of converting x from floating-point single-precision format to floating-point double-precision format
Extract	Select a field of <i>n</i> bits starting at bit position <i>b</i> in the source register, right or left justify this field in the target register, and clear all other bits of the target register to zero. This operation is used for rotate and shift instructions.
EXTS(x)	Result of extending x on the left with sign bits
GPR(x)	General-purpose register x
if...then...else...	Conditional execution, indenting shows range, else is optional
Insert	Select a field of <i>n</i> bits in the source register, insert this field starting at bit position <i>b</i> of the target register, and leave other bits of the target register unchanged. (No simplified mnemonic is provided for insertion of a field when operating on double words; such an insertion requires more than one instruction.) This operation is used for rotate and shift instructions. (Note that simplified mnemonics are referred to as extended mnemonics in the architecture specification.)
Leave	Leave innermost do loop, or the do loop described in leave statement.
MASK(x, y)	Mask having ones in positions x through y (wrapping if $x > y$ ) and zeros elsewhere.
MEM(x, y)	Contents of y bytes of memory starting at address x.
NIA	Next instruction address, which is the 32-bit address of the next instruction to be executed (the branch destination) after a successful branch. In pseudocode, a successful branch is indicated by assigning a value to NIA. For instructions which do not branch, the next instruction address is CIA + 4. Does not correspond to any architected register.
OEA	PowerPC operating environment architecture
Rotate	Rotate the contents of a register right or left <i>n</i> bits without masking. This operation is used for rotate and shift instructions.
ROTL[64](x, y)	Result of rotating the 64-bit value x left y positions
ROTL[32](x, y)	Result of rotating the 64-bit value x    x left y positions, where x is 32 bits long
Set	Bits are set to 1.
Shift	Shift the contents of a register right or left <i>n</i> bits, clearing vacated bits (logical shift). This operation is used for rotate and shift instructions.
SINGLE(x)	Result of converting x from floating-point double-precision format to floating-point single-precision format.
SPR(x)	Special-purpose register x
TRAP	Invoke the system trap handler.
Undefined	An undefined value. The value may vary from one implementation to another, and from one execution to another on the same implementation.

Table 6-2. Notation and Conventions (continued)

Notation/Convention	Meaning
UISA	PowerPC user instruction set architecture
VEA	PowerPC virtual environment architecture

Table 6-3 describes instruction field notation conventions used throughout this chapter.

Table 6-3. Instruction Field Conventions

The PowerPC Architecture Specification	Equivalent in AltiVec Technology PEM as:
RA, RB, RT, RS	rA, rB, rD, rS
SI	SIMM
U	IMM
UI	UIMM
VA, VB, VC, VT, VS	vA, vB, vC, vD, vS
<i>I, II, III</i>	0...0 (shaded)

Precedence rules for pseudocode operators are summarized in Table 6-4.

Table 6-4. Precedence Rules

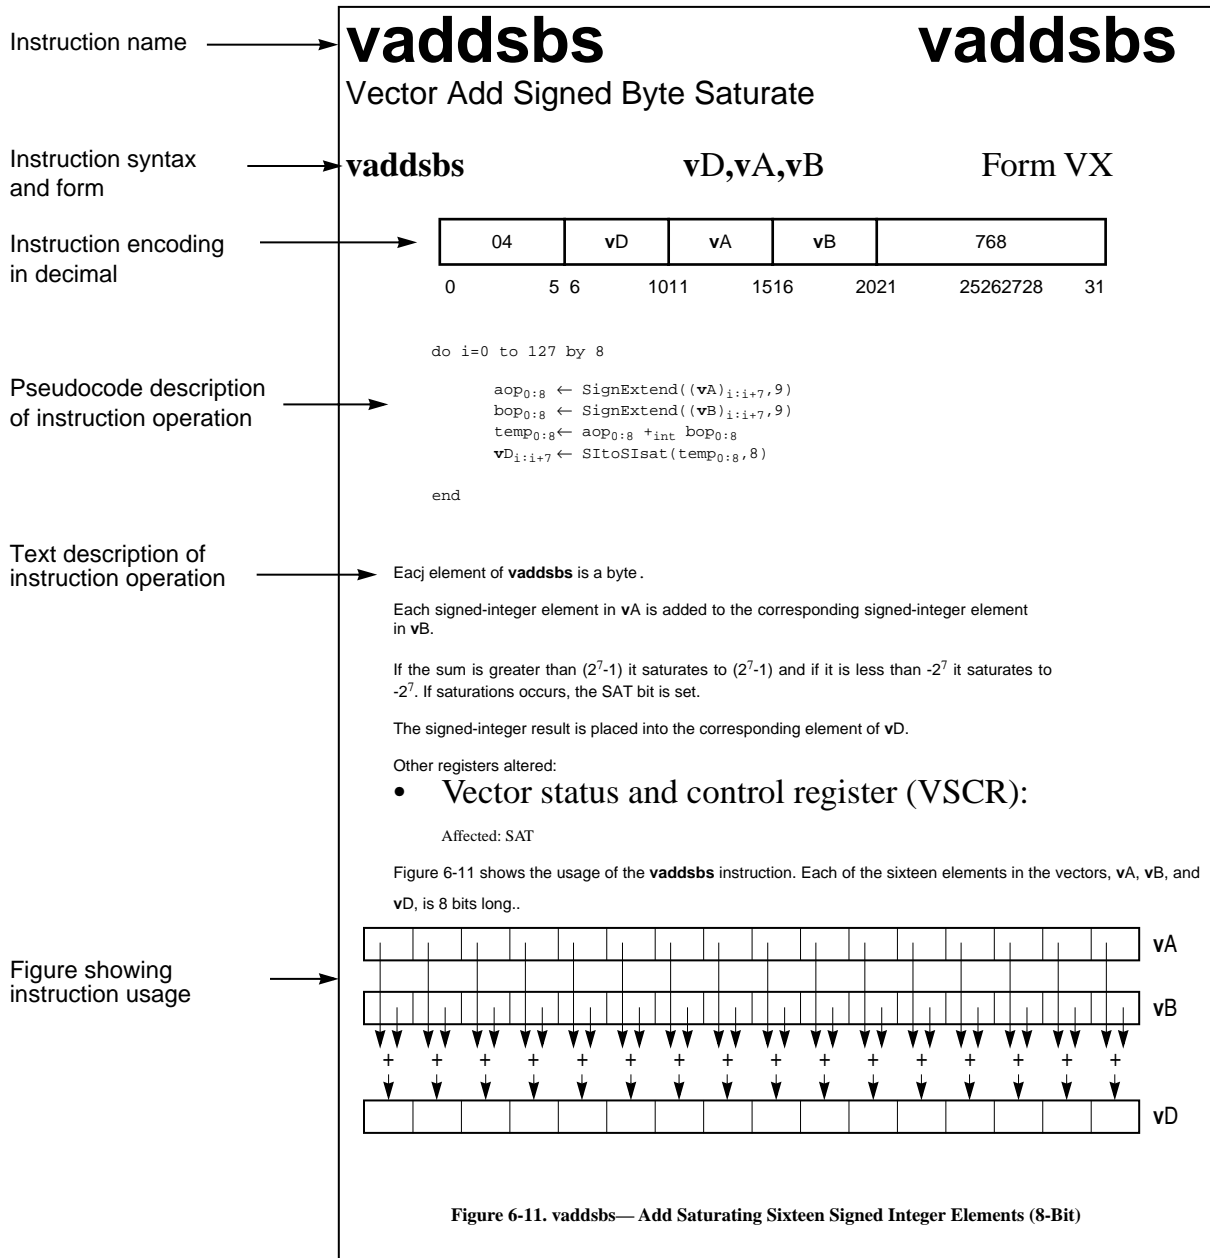
Operators	Associativity
$x[n]$ , function evaluation	Left to right
$(n)x$ or replication, $x(n)$ or exponentiation	Right to left
unary $-$ , $\neg$	Right to left
$*$ , $\div$	Left to right
$+$ , $-$	Left to right
$\parallel$	Left to right
$=$ , $\neq$ , $<$ , $\leq$ , $>$ , $\geq$ , $<U$ , $>U$ , $?$	Left to right
$\&$ , $\oplus$ , $\equiv$	Left to right
$ $	Left to right
$-$ (range), $:$ (range)	None
$\leftarrow$ , $\leftarrow_{iea}$	None

Operators higher in Table 6-4 are applied before those lower in the table. Operators at the same level in the table associate from left to right, from right to left, or not at all, as shown in the Associativity column. For example, ' $-$ ' (unary minus) associates from left to right, so  $a - b - c = (a - b) - c$ . Parentheses are used to override the evaluation order implied by

Table 6-4, or to increase clarity; parenthesized expressions are evaluated before serving as operands.

## 6.2 AltiVec Instruction Set

The remainder of this chapter lists and describes the instruction set for the AltiVec architecture. The instructions are listed in alphabetical order by mnemonic. The diagram below shows the format for each instruction description page.






# dst

# dst

Data Stream Touch


**dst**                      **rA,rB,STRM**                      (**T=0**)                      Form X  
**dstt**                      **rA,rB,STRM**                      (**T=1**)



```
addr0:63 ← (rA)
DataStreamPrefetchControl ← "start" || STRM || T || (rB) || addr
```

This instruction initiates a software directed cache prefetch. The instruction is a hint to hardware that performance will probably be improved if the cache blocks containing the specified data stream are fetched into the data cache because the program will probably soon load from the stream.

The instruction associates the data stream specified by the contents of **rA** and **rB** with the stream ID specified by **STRM**. The instruction defines a data stream **STRM** as starting at an effective address (**rA**) and having count units of size quad words separated by stride bytes (as specified in **rB**). The **T** bit of the instruction indicates whether the data stream is likely to be loaded from fairly frequently in the near future (**T = 0**) or to be transient and referenced very few times (**T = 1**).




The **dst** instruction does the following:

- Defines the characteristics of a data stream **STRM** by the contents of **rA** and **rB**
- Associates the stream with a specified stream ID, **STRM** (Range for **STRM** is 0-3)
- Indicates that the data in the specified stream **STRM** starting at the address in **rA** may soon be loaded
- Indicates whether memory locations within the stream are likely to be needed over a longer period of time (**T=0**) or be treated as transient data (**T=1**)
- Terminates prefetching from any stream that was previously associated with the specified stream ID, **STRM**.



The specified data stream is encoded for 32-bit follows:

- Effective address:  $rA$ , where  $rA \neq 0$
- Block size:  $rB[3-7]$  if  $rB[3-7] \neq 0$ ; otherwise 32
- Block count:  $rB[8-15]$  if  $rB[8-15] \neq 0$ ; otherwise 256
- Block stride:  $rB[16-31]$  if  $rB[16-31] \neq 0$ ; otherwise 32768



Other registers altered:

- None

Simplified mnemonics:

**dst**  $rA,rB,STRM$  equivalent to **dst**  $rA,rB,STRM,0$

**dstt**  $rA,rB,STRM$  equivalent to **dst**  $rA,rB,STRM,1$


For more information on the **dst** instruction, refer to Chapter 5, “Cache, Exceptions, and Memory Management.”

# dstst

# dstst

Data Stream Touch for Store

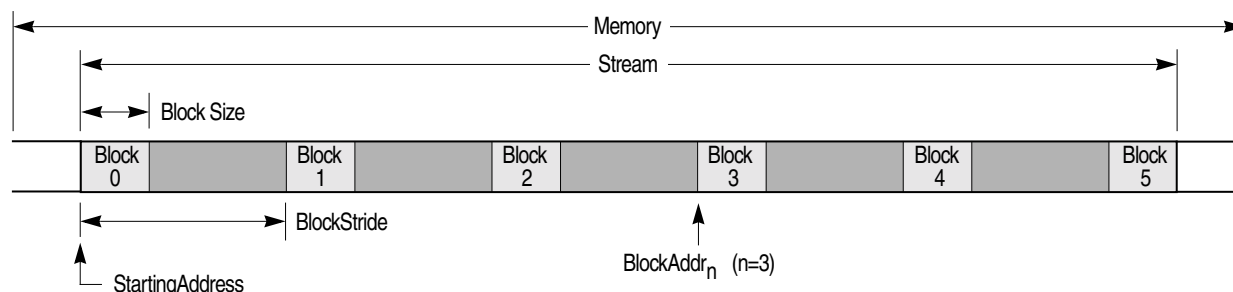
**dstst**                      **rA,rB,STRM**                      (**T=0**)                      Form X  
**dststt**                      **rA,rB,STRM**                      (**T=1**)



```
addr0:63 ← (rA)
DataStreamPrefetchControl ← "start" || T || static || (rB) || addr
```

This instruction initiates a software directed cache prefetch. The instruction is a hint to hardware that performance will probably be improved if the cache blocks containing the specified data stream are fetched into the data cache because the program will probably soon write to (store into) the stream.

The instruction associates the data stream specified by the contents of **rA** and **rB** with the stream ID specified by **STRM**. The instruction defines a data stream **STRM** as starting at an effective address (**rA**) and having count units of size quad words separated by stride bytes (as specified in **rB**). The **T** bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (**T = 0**) or to be transient and referenced very few times (**T = 1**).




The **dstst** instruction does the following:

- Defines the characteristics of a data stream **STRM** by the contents of **rA** and **rB**
- Associates the stream with a specified stream ID, **STRM** (Range for **STRM** is 0-3)
- Indicates that the data in the specified stream **STRM** starting at the address in **rA** may soon be stored in to memory
- Indicates whether memory locations within the stream are likely to be stored into fairly frequently in the near future (**T=0**) or be treated as transient data (**T=1**)
- Terminates prefetching from any stream that was previously associated with the specified stream ID, **STRM**.

The specified data stream is encoded for 32-bit follows:

- Effective address:  $rA$ , where  $rA \neq 0$
- Block size:  $rB[3-7]$  if  $rB[3-7] \neq 0$ ; otherwise 32
- Block count:  $rB[8-15]$  if  $rB[8-15] \neq 0$ ; otherwise 256
- Block stride:  $rB[16-31]$  if  $rB[16-31] \neq 0$ ; otherwise 32768



**Figure 6-1. Format of  $rB$  in  $dst$  instruction (32-bit)**

Other registers altered:

- None

Simplified mnemonics:

**dstst**  $rA,rB,STRM$  equivalent to **dstst**  $rA,rB,STRM,0$

**dststt**  $rA,rB,STRM$  equivalent to **dstst**  $rA,rB,STRM,1$

For more information on the **dstst** instruction, refer to Chapter 5, “Cache, Exceptions, and Memory Management.”

# Ivebx


Load Vector Element Byte Indexed

# Ivebx

Ivebx

vD,rA,rB

Form X



- For 32-bit:


```

if rA=0 then b ← 0
else      b ← (rA)
EA ← b + (rB)
eb ← EA28:31
vD ← undefined
if the processor is in big-endian mode
then vDeb*8:(eb*8)+7 ← MEM(EA,1)
else vD120-(eb*8):127-(eb*8) ← MEM(EA,1)
— EA = (rA|0)+(rB); m = EA[28-31] (the offset of the byte in its aligned
quadword).
```

For big-endian mode, the byte addressed by EA is loaded into byte m of vD. In little-endian mode, it is loaded into byte (15–m) of vD. Remaining bytes in vD are undefined.

Other registers altered:

- None



Note: In vector registers, x means boundedly undefined after a load and don't care after a store. In memory, x means don't care after a load, and leave at current value after a store.

Figure 6-2. Effects of Example Load/Store Instructions

Freescale Semiconductor, Inc.

# Ivehx


Load Vector Element Half Word Indexed

# Ivehx

**Ivehx**

**vD,rA,rB**

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & (~1)
eb ← EA28:31
vD ← undefined
if the processor is in big-endian mode
then vD(eb*8):(eb*8)+15 ← MEM(EA, 2)
else vD112-(eb*8):127-(eb*8) ← MEM(EA, 2)
  
```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with ~1. Let m = EA[28-30]; m is the half-word offset of the half-word in its aligned quadword in memory.

If the processor is in big-endian mode, the half-word addressed by EA is loaded into half-word m of vD. If the processor is in little-endian mode, the half-word addressed by EA is loaded into half-word (7-m) of vD. The remaining half-words s in vD are set to undefined values. Figure 6-2 shows this instruction.

Other registers altered:

- None

# lviewx


Load Vector Element Word Indexed

# lviewx

lviewx

vD,rA,rB

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & (~3)
eb ← EA28:31
vD ← undefined
if the processor is in big-endian mode
then vDeb*8:(eb*8)+31 ← MEM(EA, 4)
else vD96-(eb*8):127-(eb*8) ← MEM(EA, 4)

```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with ~3. Let m = EA[28–29]; m is the word offset of the word in its aligned quadword in memory.

If the processor is in big-endian mode, the word addressed by EA is loaded into word m of vD. If the processor is in little-endian mode, the word addressed by EA is loaded into word (3-m) of vD. The remaining words in vD are set to undefined values. Figure 6-2 shows this instruction.


Other registers altered:

- None





The above **lvsl** instruction followed by a Vector Permute (**vperm**) would do a simulated alignment of a four-element floating-point vector misaligned on quad-word boundary at address 0x0...C.



**Figure 6-4. Instruction vperm Used in Aligning Data**

Refer, also, to the description of the **lvsl** instruction for suggested uses of the **lvsl** instruction.




**lvx**

Load Vector Indexed

**lvx****lvx****vD,rA,rB**

(LRU = 0)

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & (~0xF)
if the processor is in big-endian mode
then vD ← MEM(EA,16)
else vD ← MEM(EA+8,8) || MEM(EA,8)

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with ~0xF.

If the processor is in big-endian mode, the quadword in memory addressed by EA is loaded into vD.

If the processor is in little-endian mode, the doubleword addressed by EA is loaded into vD[64–127] and the doubleword addressed by EA+8 is loaded into vD[0–63]. Note that normal little-endian PowerPC address swizzling is also performed. See Section 3.1, “Data Organization in Memory,” for more information.

Figure 6-3 shows this instruction.

Other registers altered:


- None

# lvxl

# lvxl

Load Vector Indexed LRU

**lvxl**                                      **vD,rA,rB**                                      (LRU = 1)                                      Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & (~0xF)
if the processor is in big-endian mode
then vD ← MEM(EA,16)
else vD ← MEM(EA+8,8) || MEM(EA,8)
    
```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with ~0xF.

If the processor is in big-endian mode, the quadword addressed by EA is loaded into vD.

If the processor is in little-endian mode, the doubleword addressed by EA is loaded into vD[64–127] and the doubleword addressed by EA+8 is loaded into vD[0–63]. Note that normal little-endian PowerPC address swizzling is also performed. See Section 3.1, “Data Organization in Memory,” for more information.

**lvxl** provides a hint that the program may not need quadword addressed by EA again soon.

Note that on some implementations, the hint provided by the **lvxl** instruction and the corresponding hint provided by the Store Vector Indexed LRU (**stvxli**) instruction (see Section 5.2.1.2, “Transient Streams”) are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference. Figure 6-3 shows this instruction.

Other registers altered:

- None


# mfvscr

Move from Vector Status and Control Register

# mfvscr

**mfvscr****vD**

Form VX



$$vD \leftarrow {}^9_0 || (VSCR)$$

The contents of the VSCR are placed into **vD**.

Note that the programmer should assume that **mtvscr** and **mfvscr** take substantially longer to execute than other VX instructions

Other registers altered:

- None

# mtvscr


Move to Vector Status and Control Register

# mtvscr

mtvscr

vB

Form VX



$$VSCR \leftarrow (vB)_{96:127}$$

The contents of vB are placed into the VSCR.

Other registers altered:

- None

# stvebx


Store Vector Element Byte Indexed

# stvebx

stvebx

vS,rA,rB

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← b + (rB)
eb ← EA28:31
if the processor is in big-endian mode
then MEM(EA,1) ← (vS)eb*8:(eb*8)+7
else MEM(EA,1) ← (vS)120-(eb*8):127-eb*8

```

— Let the EA be the sum (rA|0)+(rB). Let  $m = EA[28-31]$ ;  $m$  is the byte offset of the byte in its aligned quadword in memory.

If the processor is in big-endian mode, byte  $m$  of  $vS$  is stored into the byte in memory addressed by EA. If the processor is in little-endian mode, byte  $(15-m)$  of  $vS$  is stored into the byte addressed by EA. Figure 6-2 shows how a store instruction is performed for a vector register.

Other registers altered:

- None

# stvehx


Store Vector Element Half Word Indexed

# stvehx

**stvehx**

**vS,rA,rB**

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & (~0x1)
eb ← EA28:31
if the processor is in big-endian mode
then MEM(EA,2) ← (vS)eb*8:(eb*8)+15
else MEM(EA,2) ← (vS)112-eb*8:127-(eb*8)
    
```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with ~0x1. Let m = EA[28–30]; m is the half-word offset of the half-word in its aligned quadword in memory.

If the processor is in big-endian mode, half-word m of vS is stored into the half-word addressed by EA. If the processor is in little-endian mode, half-word (7-m) of vS is stored into the half-word addressed by EA. Figure 6-2 shows how a store instruction is performed for a vector register.

Other registers altered:

- None



# stvevx


Store Vector Element Word Indexed

# stvevx

stvevx

vS,rA,rB

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFC
eb ← EA28:31
if the processor is in big-endian mode
then MEM(EA,4) ← (vS)eb*8:(eb*8)+31
else MEM(EA,4) ← (vS)96-eb*8:127-(eb*8)

```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFC.

Let m = EA[28-29]; m is the word offset of the word in its aligned quadword in memory.

If the processor is in big-endian mode, word m of vS is stored into the word addressed by EA. If the processor is in little-endian mode, word (3-m) of vS is stored into the word addressed by EA. Figure 6-2 shows how a store instruction is performed for a vector register.

Other registers altered:


- None

# stvx

Store Vector Indexed

# stvx

**stvx**                                      **vS,rA,rB**                                      (LRU = 0)                                      Form X



- For 32-bit:

```

if rA=0 then b ← 0
else b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFF0
if the processor is in big-endian mode
then MEM(EA,16) ← (vS)
else MEM(EA,16) ← (vS)64:127 || (vS)0:63
    
```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFF0.

If the processor is in big-endian mode, the contents of vS are stored into the quadword addressed by EA. If the processor is in little-endian mode, the contents of vS[64–127] are stored into the doubleword addressed by EA, and the contents of vS[0–63] are stored into the doubleword addressed by EA+8.

**stvxl** and **stvxlt** provide a hint that the quadword addressed by EA will probably not be needed again by the program in the near future.

Figure 6-2 shows how a store instruction is performed for a vector register.

Other registers altered:

- None

# stvxl


Store Vector Indexed LRU

# stvxl

**stvxl****vS,rA,rB**

(LRU = 1)

Form X



- For 32-bit:

```

if rA=0 then b ← 0
else b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFF0
if the processor is in big-endian mode
then MEM(EA,16) ← (vS)
else MEM(EA,16) ← (vS)64:127 || (vS)0:63

```

— Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFF0.

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFF0. If the processor is in big-endian mode, the contents of vS are stored into the quadword addressed by EA. If the processor is in little-endian mode, the contents of vS[64–127] are stored into the doubleword addressed by EA, and the contents of vS[0–63] are stored into the doubleword addressed by EA+8. The **stvxl** and **stvxlt** instructions provide a hint that the quad word addressed by EA will probably not be needed again by the program in the near future.

Note that on some implementations, the hint provided by the **stvxl** instruction (see Section 5.2.2, “Prioritizing Cache Block Replacement”) is applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference. Figure 6-2 shows how a store instruction is performed on the vector registers.

Other registers altered:

- None


# vaddcuw

Vector Add Carryout Unsigned Word

# vaddcuw

**vaddcuw**                      **vD,vA,vB**

Form VX




```
do i=0 to 127 by 32
    aop0:32 ← ZeroExtend((vA)i:i+31, 33)
    bop0:32 ← ZeroExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int bop0:32
    vDi:i+31 ← ZeroExtend(temp0, 32)
end
```

Each unsigned-integer word element in **vA** is added to the corresponding unsigned-integer word element in **vB**. The carry out of bit 0 of the 32-bit sum is zero-extended to 32 bits and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-5 shows the usage of the **vaddcuw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-5. vaddcuw—Determine Carries of Four Unsigned Integer Adds (32-Bit)**


# vaddfp

Vector Add Floating Point

# vaddfp

**vaddfp****vD,vA,vB**

Form VX



```

do i = 0,127,32
    (vD)i:i+31 ← RndToNearFP32((vA)i:i+31 +fp (vB)i:i+31)
end

```


The four 32-bit floating-point values in **vA** are added to the four 32-bit floating-point values in **vB**. The four intermediate results are rounded and placed in **VD**.

If **VSCR[NJ] = 1**, every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign.

Other registers altered:

- None

Figure 6-6 shows the usage of the **vaddfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-6. vaddfp—Add Four Floating-Point Elements (32-Bit)**


# vaddsbs

Vector Add Signed Byte Saturate

# vaddsbs

**vaddsbs**                      **vD,vA,vB**

Form VX



```
do i=0 to 127 by 8
    aop0:8 ← SignExtend((vA)i:i+7, 9)
    bop0:8 ← SignExtend((vB)i:i+7, 9)
    temp0:8 ← aop0:8 +int bop0:8
    vDi:i+7 ← SItoSIsat(temp0:8, 8)
end
```

Each element of **vaddsbs** is a byte.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the sum is greater than  $(2^7-1)$  it saturates to  $(2^7-1)$  and if it is less than  $-2^7$  it saturates to  $-2^7$ . If saturation occurs, the SAT bit is set.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
 Affected: SAT

Figure 6-7 shows the usage of the **vaddsbs** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-7. vaddsbs—Add Saturating Sixteen Signed Integer Elements (8-Bit)**


# vaddshs

Vector Add Signed Half Word Saturate

# vaddshs

**vaddshs**                      **vD,vA,vB**

Form VX



```
do i=0 to 127 by 16
    aop0:16 ← SignExtend((vA)i:i+15,16)
    bop0:16 ← SignExtend((vB)i:i+15,16)
    temp0:16 ← aop0:16 +int bop0:16
    vDi:i+15 ← SItoSIsat(temp0:16,16)
end
```

Each element of **vaddshs** is a half word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the sum is greater than  $(2^{15}-1)$  it saturates to  $(2^{15}-1)$  and if it is less than  $-2^{15}$  it saturates to  $-2^{15}$ . If saturation occurs, the SAT bit is set.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
Affected: SAT

Figure 6-8 shows the usage of the **vaddshs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-8. vaddshs— Add Saturating Eight Signed Integer Elements (16-Bit)**


# vaddsws

Vector Add Signed Word Saturate

# vaddsws

**vaddsws**                      **vD,vA,vB**

Form VX



```
do i=0 to 127 by 32
    aop0:32 ← SignExtend((vA)i:i+31, 33)
    bop0:32 ← SignExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int bop0:32
    vDi:i+31 ← SItoSIsat(temp0:32, 32)
end
```

Each element of **vaddsws** is a word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the sum is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $(-2^{31})$  it saturates to  $(-2^{31})$ . If saturation occurs, the SAT bit is set.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
 Affected: SAT

Figure 6-9 shows the usage of the **vaddsws** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-9. vaddsws—Add Saturating Four Signed Integer Elements (32-Bit)**



# vaddubm


Vector Add Unsigned Byte Modulo

# vaddubm

**vaddubm**

**vD,vA,vB**

Form VX



```
do i=0 to 127 by 8
    vDi:i+7 ← (vA)i:i+7 +int (vB)i:i+7
end
```

Each element of **vaddubm** is a byte.


Each integer element in **vA** is modulo added to the corresponding integer element in **vB**. The integer result is placed into the corresponding element of **vD**.

Note that the **vaddubm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None

Figure 6-10 shows the **vaddubm** instruction usage. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-10. vaddubm—Add Sixteen Integer Elements (8-Bit)**


# vaddubs

Vector Add Unsigned Byte Saturate

# vaddubs

**vaddubs**                      **vD,vA,vB**

Form VX



```

do i=0 to 127 by 8
    aop0:8 ← ZeroExtend((vA)i:i+7, 9)
    bop0:8 ← ZeroExtend((vB)i:i+7, 9)
    temp0:8 ← aop0:8 +int bop0:8
    vDi:i+7 ← UItoUISat(temp0:8, 8)
end
  
```

Each element of **vaddubs** is a byte.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.


If the sum is greater than  $(2^8-1)$  it saturates to  $(2^8-1)$  and the SAT bit is set.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
   Affected: SAT

Figure 6-11 shows the usage of the **vaddubs** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-11. vaddubs—Add Saturating Sixteen Unsigned Integer Elements (8-Bit)**

# vadduhm


Vector Add Unsigned Half Word Modulo

# vadduhm

**vadduhm**

**vD,vA,vB**

Form VX



```
do i=0 to 127 by 16
    vDi:i+15 ← (vA)i:i+15 +int (vB)i:i+15
end
```

Each element of **vadduhm** is a half word.


Each integer element in **vA** is added to the corresponding integer element in **vB**. The integer result is placed into the corresponding element of **vD**.

Note that the **vadduhm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None

Figure 6-12 shows the usage of the **vadduhm** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-12. vadduhm—Add Eight Integer Elements (16-Bit)**


# vadduhs

Vector Add Unsigned Half Word Saturate

# vadduhs

**vadduhs**                      **vD,vA,vB**

Form VX



```
do i=0 to 127 by 16
    aop0:16 ← ZeroExtend((vA)i:i+15,17)
    bop0:16 ← ZeroExtend((vB)i:i+15,17)
    temp0:16 ← aop0:16 +int bop0:16
    vDi:i+15 ← UItoUISat(temp0:16,16)
end
```

Each element of **vadduhs** is a half word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.


If the sum is greater than  $(2^{16}-1)$  it saturates to  $(2^{16}-1)$  and the SAT bit is set.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
 Affected: SAT

Figure 6-13 shows the usage of the **vadduhs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-13. vadduhs—Add Saturating Eight Unsigned Integer Elements (16-Bit)**

# vadduwm


Vector Add Unsigned Word Modulo

# vadduwm

**vadduwm**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← (vA)i:i+31 +int (vB)i:i+31
end
```

Each element of **vadduwm** is a word.

Each integer element in **vA** is modulo added to the corresponding integer element in **vB**. The integer result is placed into the corresponding element of **vD**.

Note that the **vadduwm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None

Form:

- VX

Figure 6-14 shows the usage of the **vadduwm** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.




Figure 6-14. **vadduwm**—Add Four Integer Elements (32-Bit)


# vadduws

Vector Add Unsigned Word Saturate

# vadduws

**vadduws**                      **vD,vA,vB**

Form: VX



```

do i=0 to 127 by 3
    aop0:32 ← ZeroExtend((vA)i:i+31, 33)
    bop0:32 ← ZeroExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int bop0:32
    vDi:i+31 ← UItoUISat(temp0:32, 32)
end
    
```

Each element of **vadduws** is a word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.


If the sum is greater than  $(2^{32}-1)$  it saturates to  $(2^{32}-1)$  and the SAT bit is set.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
 Affected: SAT

Figure 6-15 shows the usage of the **vadduws** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-15. vadduws—Add Saturating Four Unsigned Integer Elements (32-Bit)**

# vand


Vector Logical AND

# vand

**vand**

**vD,vA,vB**

Form: VX




$$vD \leftarrow (vA) \& (vB)$$

The contents of **vA** are bitwise ANDed with the contents of **vB** and the result is placed into **vD**.

Other registers altered:

- None

Figure 6-16 shows usage of the **vand** instruction.



**Figure 6-16. vand—Logical Bitwise AND**

# vandc


Vector Logical AND with Complement

# vandc

**vandc**

**vD,vA,vB**

Form: VX




$$vD \leftarrow (vA) \& \neg(vB)$$

The contents of **vA** are ANDed with the one's complement of the contents of **vB** and the result is placed into **vD**.

Other registers altered:

- None

Figure 6-16 shows usage of the **vandc** instruction.



**Figure 6-17. vand—Logical Bitwise AND with Complement**




# vavg**sb**

Vector Average Signed Byte

# vavg**sb**

**vavg**sb**** vD,vA,vB

Form: VX



```

do i=0 to 127 by 8
    aop0:8 ← SignExtend((vA)i:i+7,9)
    bop0:8 ← SignExtend((vB)i:i+7,9)
    temp0:8 ← aop0:8 +int bop0:8 +int 1
    vDi:i+7 ← temp0:7
end
    
```


Each element of **vavg**sb**** is a byte.

Each signed-integer byte element in vA is added to the corresponding signed-integer byte element in vB, producing an 9-Bit signed-integer sum. The sum is incremented by 1. The high-order 8 bits of the result are placed into the corresponding element of vD.

Other registers altered:

- None

Figure 6-18 shows the usage of the **vavg**sb**** instruction. Each of the sixteen elements in the vectors, vA, vB, and vD, is 8 bits long.



**Figure 6-18. vavg**sb**— Average Sixteen Signed Integer Elements (8-Bit)**

# vavgsh


Vector Average Signed Half Word

# vavgsh

**vavgsh**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
  aop0:16 ← SignExtend((vA)i:i+15,17)
  bop0:16 ← SignExtend((vB)i:i+15,17)
  temp0:16 ← aop0:15 +int bop0:15 +int 1
  vDi:i+15 ← temp0:15
end
```


Each element of **vavgsh** is a half word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**, producing an 17-bit signed-integer sum. The sum is incremented by 1. The high-order 16 bits of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-19 shows the usage of the **vavgsh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-19. vavgsh—Average Eight Signed Integer Elements (16-bits)**

# vavgsw


Vector Average Signed Word

# vavgsw

**vavgsw**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 32
    aop0:32 ← SignExtend((vA)i:i+31, 33)
    bop0:32 ← SignExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int bop0:32 +int 1
    vDi:i+31 ← temp0:31
end
    
```


Each element of **vavgsw** is a word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**, producing an 33-bit signed-integer sum. The sum is incremented by 1. The high-order 32 bits of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-20 shows the usage of the **vavgsw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-20. vavgsw— Average Four Signed Integer Elements (32-Bit)**


# vavgub

Vector Average Unsigned Byte

# vavgub

**vavgub**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    aop0:8 ← ZeroExtend((vA)i:i+7, 9)
    bop0:n ← ZeroExtend((vB)i:i+7, 9)
    temp0:n ← aop0:8 +int bop0:8 +int 1
    vDi:i+7 ← temp0:7
end
```

Each element of **vavgub** is a byte.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing an 9-bit unsigned-integer sum. The sum is incremented by 1. The high-order 8 bits of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-21 shows the usage of the **vavgub** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.




Figure 6-21. vavgub—Average Sixteen Unsigned Integer Elements (8-bits)


# vavguh

# vavguh

Vector Average Unsigned Half Word

**vavguh**  $vD, vA, vB$

Form: VX



```

do i=0 to 127 by 16
    aop0:16 ← ZeroExtend((vA)i:i+15, 17)
    bop0:16 ← ZeroExtend((vB)i:i+15, 17)
    temp0:16 ← aop0:16 +int bop0:16 +int 1
    vDi:i+15 ← temp0:15
end
    
```

Each element of **vavguh** is a half word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing a 17-bit unsigned-integer. The sum is incremented by 1. The high-order 16 bits of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-22 shows the usage of the **vavgsh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.




Figure 6-22. vavgsh— Average Eight Signed Integer Elements (16-Bit)


# vavguw

Vector Average Unsigned Word

# vavguw

**vavguw**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    aop0:32 ← ZeroExtend((vA)i:i+31, 33)
    bop0:32 ← ZeroExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int bop0:32 +int 1
    vDi:i+31 ← temp0:31
end
```


Each element of **vavguw** is a word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing an 33-bit unsigned-integer sum. The sum is incremented by 1. The high-order 32 bits of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-23 shows the usage of the **vavguw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-23. vavguw—Average Four Unsigned Integer Elements (32-Bit)**

# vcfsx


Vector Convert from Signed Fixed-Point Word

# vcfsx

vcfsx

vD,vB,UIMM

Form: VX



```
do i=0 to 127 by 32
```

$$vD_{i:i+31} \leftarrow \text{CnvtSI32ToFP32}((vB)_{i:i+31}) \div_{fp} 2^{UIMM}$$


```
end
```

Each signed fixed-point integer word element in **vB** is converted to the nearest single-precision floating-point value. The result is divided by  $2^{UIMM}$  (UIMM = Unsigned immediate value) and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-24 shows the usage of the **vcfsx** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-24. vcfsx—Convert Four Signed Integer Elements to Four Floating-Point Elements (32-Bit)**


# vcfux

# vcfux

Vector Convert from Unsigned Fixed-Point Word

**vcfux**                      **vD,vB,UIMM**

Form: VX




```
do i=0 to 127 by 32
    vDi:i+31 ← CnvtUI32ToFP32((vB)i:i+31) ÷fp 2UIMM
end
```

Each unsigned fixed-point integer word element in **vB** is converted to the nearest single-precision floating-point value. The result is divided by  $2^{UIMM}$  and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-25 shows the usage of the **vcfux** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-25. vcfux—Convert Four Unsigned Integer Elements to Four Floating-Point Elements (32-Bit)**




**vcmpbfp**

Vector Compare Bounds Floating Point

**vcmpbfp**

**vcmpbfp**                      **vD,vA,vB**                      (**Rc = 0**)                      Form: VXR  
**vcmpbfp.**                      **vD,vA,vB**                      (**Rc = 1**)



```

do i=0 to 127 by 32
    le ← ((vA)i:i+31 ≤fp (vB)i:i+31)
    ge ← ((vA)i:i+31 ≥fp -(vB)i:i+31)
    vDi:i+31 ← -le || -ge || 300
end
if Rc=1 then do
    ib ← (vD = 1280)
    CR24:27 ← 0b00 || ib || 0b0
end

```

Each single-precision word element in **vA** is compared to the corresponding element in **vB**. A 2-bit value is formed that indicates whether the element in **vA** is within the bounds specified by the element in **vB**, as follows.

Bit 0 of the 2-bit value is zero if the element in **vA** is less than or equal to the element in **vB**, and is one otherwise. Bit 1 of the 2-bit value is zero if the element in **vA** is greater than or equal to the negative of the element in **vB**, and is one otherwise.

The 2-bit value is placed into the high-order two bits of the corresponding word element (bits 0–1 for word element 0, bits 32–33 for word element 1, bits 64–65 for word element 2, bits 96–97 for word element 3) of **vD** and the remaining bits of the element are cleared.

If **Rc=1**, **CR** Field 6 is set to indicate whether all four elements in **vA** are within the bounds specified by the corresponding element in **vB**, as follows.

- **CR6** = 0b00 || all\_within\_bounds || 0

Note that if any single-precision floating-point word element in **vB** is negative; the corresponding element in **vA** is out of bounds. Note that if a **vA** or a **vB** element is a NaN, the two high order bits of the corresponding result will both have the value 1.

If **VSCR[NJ] = 1**, every denormalized operand element is truncated to 0 before the comparison is made.

Other registers altered:

- Condition register (**CR6**):  
Affected: Bit 2                      (if **Rc = 1**)

Figure 6-26 shows the usage of the `vcmpbfp` instruction. Each of the four elements in the vectors, `vA`, `vB`, and `vD`, is 32 bits long.




Figure 6-26. `vcmpbfp`—Compare Bounds of Four Floating-Point Elements (32-Bit)


# vcmpeqfp

Vector Compare Equal-to-Floating Point

# vcmpeqfp

**vcmpeqfp**  
**vcmpeqfp.****vD,vA,vB**  
**vD,vA,vB**

Form: VXR



```

do i=0 to 127 by 32
  if (vA)i:i+31 =fp (vB)i:i+31
    then vDi:i+31 ← 0xFFFF_FFFF
    else vDi:i+31 ← 0x0000_0000
end
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR24:27 ← t || 0b0 || f || 0b0
end

```

Each single-precision floating-point word element in **vA** is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is equal to the element in **vB**, and is cleared to all 0s otherwise.

If  $Rc = 1$ . CR6 filed is set according to all, some, or none of the elements pairs compare equal.


- CR6 = all\_equal || 0b0 || none\_equal || 0b0

Note that if a **vA** or **vB** element is a NaN, the corresponding result will be 0x0000\_0000.

Other registers altered:

- Condition register (CR6):  
Affected: Bits 0-3 (if  $Rc = 1$ )

Figure 6-27 shows the usage of the **vcmpeqfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-27. vcmpeqfp—Compare Equal of Four Floating-Point Elements (32-Bit)**








# vcmpgefp

# vcmpgefp

Vector Compare Greater-Than-or-Equal-to Floating Point

**vcmpgefp**                      **vD,vA,vB**                      (**Rc = 0**)                      Form: VXR  
**vcmpgefp.**                      **vD,vA,vB**                      (**Rc = 1**)



```
do i=0 to 127 by 32
  if (vA)i:i+31 ≥fp (vB)i:i+31
  then vDi:i+31 ← 0xFFFF_FFFF
  else vDi:i+31 ← 0x0000_0000
end
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR24:27 ← t || 0b0 || f || 0b0
end
```

Each single-precision floating-point word element in **vA** is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is greater than or equal to the element in **vB**, and is cleared to all 0s otherwise.

If **Rc = 1**, **CR6** is set according to **all\_greater\_or\_equal || some\_greater\_or\_equal || none\_great\_or\_equal**.


$$CR6 = \text{all\_greater\_or\_equal} \parallel 0b0 \parallel \text{none\_greater\_or\_equal} \parallel 0b0.$$

Note that if a **vA** or **vB** element is a NaN, the corresponding results will be 0x0000\_0000.

Other registers altered:

- Condition register (**CR6**):  
 Affected: Bits 0-3                      (if **Rc = 1**)

Figure 6-31 shows the usage of the **vcmpgefp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long



**Figure 6-31. vcmpgefp—Compare Greater-Than-or-Equal of Four Floating-Point Elements (32-Bit)**


# vcmpgtfp

Vector Compare Greater-Than Floating-Point

# vcmpgtfp

**vcmpgtfp**                    **vD,vA,vB**  
**vcmpgtfp.**                    **vD,vA,vB**

Form: VXR



```
do i=0 to 127 by 32
    if (vA)i:i+31 >fp (vB)i:i+31
        then vDi:i+31 ← 0xFFFF_FFFF
        else vDi:i+31 ← 0x0000_0000
    end
    if Rc=1 then do
        t ← (vD = 1281)
        f ← (vD = 1280)
        CR[24:27] ← t || 0b0 || f || 0b0
    end
end
```

Each single-precision floating-point word element in **vA** is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

If **Rc = 1**, **CR6** is set according to **all\_greater\_than || some\_greater\_than || none\_greater\_than**.


$$CR6 = \text{all\_greater\_than} \parallel 0b0 \parallel \text{none greater\_than} \parallel 0b0.$$

Note that if a **vA** or **vB** element is a NaN, the corresponding results will be 0x0000\_0000.

Other registers altered:

- Condition register (**CR6**):  
 Affected: Bits 0-3 (if **Rc = 1**)

Figure 6-32 shows the usage of the **vcmpgtfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-32. vcmpgtfp—Compare Greater-Than of Four Floating-Point Elements (32-Bit)**




**vcmpgtsbx**

Vector Compare Greater-Than Signed Byte

**vcmpgtsbx****vcmpgtsb****vD,vA,vB**

Form: VXR

**vcmpgtsb.****vD,vA,vB**

```

do i=0 to 127 by 8
    if (vA)i:i+7 >si (vB)i:i+7
        then vDi:i+7 ← 81
        else vDi:i+7 ← 80
    end
end
if Rc=1 then do
    t ← (vD = 1281)
    f ← (vD = 1280)
    CR24:27 ← t || 0b0 || f || 0b0
end

```

Each element of **vcmpgtsb** is a byte.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

If  $Rc = 1$ , CR6 is set according to all\_greater\_than || some\_greater\_than || none\_greater\_than.

CR6 = all\_greater\_than || 0b0 || none\_greater\_than || 0b0.

Note that if a **vA** or **vB** element is a NaN, the corresponding results will be 0x0000\_0000.

Other registers altered:

- Condition register (CR6):  
Affected: Bits 0-3 (if  $Rc = 1$ )

Figure 6-33 shows the usage of the **vcmpgtsb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.




Figure 6-33. **vcmpgtsb**—Compare Greater-Than of Sixteen Signed Integer Elements (8-Bit)




# vcmpgtswx

Vector Compare Greater-Than Signed Word

# vcmpgtswx

**vcmpgtsw**                    **vD,vA,vB**  
**vcmpgtsw.**                 **vD,vA,vB**

Form: VXR



```

do i=0 to 127 by 32
  if (vA)i:i+31 >si (vB)i:i+31
    then vDi:i+31 ← 321
    else vDi:i+31 ← 320
  end
  if Rc=1 then do
    t ← (vD = 1281)
    f ← (vD = 1280)
    CR24:27 ← t || 0b0 || f || 0b0
  end
end

```

Each element of **vcmpgtsw** is a word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

If  $Rc = 1$ , CR6 is set according to `all_greater_than || some_greater_than || none_greater_than`.

CR6 = `all_greater_than || 0b0 || none_greater_than || 0b0`.

Note that if a **vA** or **vB** element is a NaN, the corresponding results will be `0x0000_0000`.

Other registers altered:

- Condition register (CR6):  
Affected: Bits 0-3 (if  $Rc = 1$ )

Figure 6-35 shows the usage of the **vcmpgtsw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.




Figure 6-35. **vcmpgtsw**—Compare Greater-Than of Four Signed Integer Elements (32-Bit)




**vcmpgtuhx**

Vector Compare Greater-Than Unsigned Half Word

**vcmpgtuhx**

**vcmpgtuh**                    **vD,vA,vB**  
**vcmpgtuh.**                 **vD,vA,vB**

Form: VXR



```

do i=0 to 127 by 16
    if (vA)i:i+15 >ui (vB)i:i+15
        then vDi:i+15 ← 161
        else vDi:i+15 ← 160
    end
    if Rc=1 then do
        t ← (vD = 1281)
        f ← (vD = 1280)
        CR[24-27] ← t || 0b0 || f || 0b0
    end
end

```

Each element of **vcmpgtuh** is a half word. Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

If  $Rc = 1$ , CR6 is set according to `all_greater_than || some_greater_than || none_greater_than`.


CR6 = `all_greater_than || 0b0 || none_greater_than || 0b0`.

Note that if a **vA** or **vB** element is a NaN, the corresponding results will be 0x0000\_0000.

Other registers altered:

- Condition register (CR6):  
Affected: Bits 0-3 (if  $Rc = 1$ )

Figure 6-37 shows the usage of the **vcmpgtuh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-37. vcmpgtuh—Compare Greater-Than of Eight Unsigned Integer Elements (16-Bit)**



# vctsxS


# vctsxS

Vector Convert to Signed Fixed-Point Word Saturate

vctsxS

vD,vB,UIMM

Form: VX



```

do i=0 to 127 by 32
  if (vB)i+1:i+8=255 | (vB)i+1:i+8 + UIMM ≤ 254 then
    vDi:i+31 ← CnvtFP32ToSI32Sat((vB)i:i+31 *fp 2UIMM)
  else
    do
      if (vB)i=0 then vDi:i+31 ← 0x7FFF_FFFF
      else vDi:i+31 ← 0x8000_0000
      VSCRSAT ← 1
    end
  end
end
    
```


Each single-precision word element in **vB** is multiplied by  $2^{UIMM}$ . The product is converted to a signed integer using the rounding mode, Round toward Zero. If the intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$ ; if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . A signed-integer result is placed into the corresponding word element of **vD**.

Fixed-point integers used by the vector convert instructions can be interpreted as consisting of 32-UIMM integer bits followed by UIMM fraction bits. The vector convert to fixed-point word instructions support only the rounding mode, Round toward Zero. A single-precision number can be converted to a fixed-point integer using any of the other three rounding modes by executing the appropriate vector round to floating-point integer instruction before the vector convert to fixed-point word instruction.

Other registers altered:

- Vector status and control register (VSCR):  
Affected: SAT

Figure 6-39 shows the usage of the **vctsxS** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-39. vctsxS—Convert Four Floating-Point Elements to Four Signed Integer Elements (32-Bit)**


# vctuxs

# vctuxs

Vector Convert to Unsigned Fixed-Point Word Saturate

**vctuxs**                      **vD,vB,UIMM**

Form: VX



```

do i=0 to 127 by 32
  if (vB)i+1:i+8=255 | (vB)i+1:i+8 + UIMM ≤ 254 then
    vDi:i+31 ← CnvtFP32ToUI32Sat((vB)i:i+31 *fp 2UIMM)
  else
    do
      if (vB)i=0 then vDi:i+31 ← 0xFFFF_FFFF
      else vDi:i+31 ← 0x0000_0000
      VSCRSAT ← 1
    end
  end
end
    
```

Each single-precision floating-point word element in **vB** is multiplied by  $2^{UIMM}$ . The product is converted to an unsigned fixed-point integer using the rounding mode Round toward Zero.


If the intermediate result is greater than  $(2^{32}-1)$  it saturates to  $(2^{32}-1)$  and if it is less than 0 it saturates to 0.

The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):  
     Affected: SAT

Figure 6-40 shows the usage of the **vctuxs** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-40. vctuxs—Convert Four Floating-Point Elements to Four Unsigned Integer Elements (32-Bit)**



# vexptefp


# vexptefp

Vector 2 Raised to the Exponent Estimate Floating Point

vexptefp

vD,vB

Form: VX



```
do i=0 to 127 by 32
    x ← (vB)i:i+31
    vDi:i+31 ← 2x
end
```

The single-precision floating-point estimate of 2 raised to the power of each single-precision floating-point element in vB is placed into the corresponding element of vD.

The estimate has a relative error in precision no greater than one part in 16, that is,

$$\left| \frac{\text{estimate} - 2^x}{2^x} \right| \leq \frac{1}{16}$$

where  $x$  is the value of the element in vB. The most significant 12 bits of the estimate's significant are monotonic. Note that the value placed into the element of vD may vary between implementations, and between different executions on the same implementation.

If an operation has an integral value and the resulting value is not 0 or  $+\infty$ , the result is exact.

Operation with various special values of the element in vB is summarized in Table 6-5 below.

**Table 6-5. Special Values of the Element in vB**


Value of Element in vB	Result
$-\infty$	+0
-0	+1
+0	+1
$+\infty$	$+\infty$
NaN	QNaN

If VSCR[NJ] = 1, every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign.

Other registers altered:

- None

Figure 6-41 shows the usage of the **vexptefp** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-41. vexptefp—2 Raised to the Exponent Estimate Floating-Point for Four Floating-Point Elements (32-Bit)**

# vlogefp


Vector Log<sub>2</sub> Estimate Floating Point

# vlogefp

**vlogefp**

**vD, vB**

Form: VX



```
do i=0 to 127 by 32
    x ← (vB)i:i+31
    vDi:i+31 ← log2(x)
end
```

The single-precision floating-point estimate of the base 2 logarithm of each single-precision floating-point element in **vB** is placed into the corresponding element of **vD**.

The estimate has an absolute error in precision (absolute value of the difference between the estimate and the infinitely precise value) no greater than 2<sup>-5</sup>. The estimate has a relative error in precision no greater than one part in 8, as described below:

$$\left( \left| \text{estimate} - \log_2(x) \right| \leq \frac{1}{32} \right) \quad \text{unless} \quad |x - 1| \leq \frac{1}{8}$$

where *x* is the value of the element in **vB**, except when  $|x-1| \leq 1 \div 8$ . The most significant 12 bits of the estimate's significant are monotonic. Note that the value placed into the element of **vD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the element in **vB** is summarized below in Table 6-6.

**Table 6-6. Special Values of the Element in vB**


Value	Result
-∞	QNaN
less than 0	QNaN
±0	-∞
+∞	+∞
NaN	QNaN

If VSCR[NJ] = 1, every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign.

Other registers altered:

- None

Figure 6-42 shows the usage of the **vexptfp** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-42. vexptfp—Log<sub>2</sub> Estimate Floating-Point for Four Floating-Point Elements (32-Bit)**


# vmaddfp

Vector Multiply Add Floating Point

# vmaddfp

**vmaddfp****vD,vA,vC,vB**

Form: VA



```

do i=0 to 127 by 32
    vDi:i+31 ← RndToNearFP32(((vA)i:i+31 *fp (vC)i:i+31) +fp (vB)i:i+31)
end

```


Each single-precision floating-point word element in **vA** is multiplied by the corresponding single-precision floating-point word element in **vC**. The corresponding single-precision floating-point word element in **vB** is added to the product. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

Note that a vector multiply floating-point instruction is not provided. The effect of such an instruction can be obtained by using **vmaddfp** with **vB** containing the value -0.0 (0x8000\_0000) in each of its four single-precision floating-point word elements. (The value must be -0.0, not +0.0, in order to obtain the IEEE-conforming result of -0.0 when the result of the multiplication is -0.)

Other registers altered:

- None

If VSCR[NJ] = 1, every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign. Figure 6-43 shows the usage of the **vmaddfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-43. vmaddfp—Multiply-Add Four Floating-Point Elements (32-Bit)**

# vmaxfp


Vector Maximum Floating Point

# vmaxfp

**vmaxfp**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    if (vA)i:i+31 ≥fp (vB)i:i+31
        then vDi:i+31 ← (vA)i:i+31
        else vDi:i+31 ← (vB)i:i+31
end
```


Each single-precision floating-point word element in **vA** is compared to the corresponding single-precision floating-point word element in **vB**. The larger of the two single-precision floating-point values is placed into the corresponding word element of **vD**.

The maximum of +0 and -0 is +0. The maximum of any value and a NaN is a QNaN.

Other registers altered:

- None

Figure 6-44 shows the usage of the **vmaxfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-44. vmaxfp—Maximum of Four Floating-Point Elements (32-Bit)**

# vmaxsb


# vmaxsb

Vector Maximum Signed Byte

**vmaxsb**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 8
    if (vA)i:i+7 ≥si (vB)i:i+7
        then vDi:i+7 ← (vA)i:i+7
        else vDi:i+7 ← (vB)i:i+7
end
    
```


Each element of **vmaxsb** is a byte.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-45 shows the usage of the **vmaxsb** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-45. vmaxsb—Maximum of Sixteen Signed Integer Elements (8-Bit)**


# vmaxsh

# vmaxsh

Vector Maximum Signed Half Word

**vmaxsh**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
  if (vA)i:i+7 ≥si (vB)i:i+15
    then vDi:i+15 ← (vA)i:i+15
    else vDi:i+15 ← (vB)i:i+15
end
```


Each element of **vmaxsh** is a half word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-46 shows the usage of the **vmaxsh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-46. vmaxsh—Maximum of Eight Signed Integer Elements (16-Bit)**



# vmaxsw


Vector Maximum Signed Word

# vmaxsw

**vmaxsw**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 32
    if (vA)i:i+31 ≥si (vB)i:i+31
        then vDi:i+31 ← (vA)i:i+31
        else vDi:i+31 ← (vB)i:i+31
end
    
```


Each element of **vmaxsw** is a word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-47 shows the usage of the **vmaxsw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-47. vmaxsw—Maximum of Four Signed Integer Elements (32-Bit)**


# vmaxub

# vmaxub

Vector Maximum Signed Byte

**vmaxub**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    if (vA)i:i+7 ≥ui (vB)i:i+7
        then vDi:i+7 ← (vA)i:i+7
        else vDi:i+7 ← (vB)i:i+7
end
```


Each element of **vmaxub** is a byte.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-48 shows the usage of the **vmaxub** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-48. vmaxub—Maximum of Sixteen Unsigned Integer Elements (8-Bit)**

# vmaxuh


# vmaxuh

Vector Maximum Unsigned Half Word

**vmaxuh**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 16
    if (vA)i:i+15 ≥ui (vB)i:i+15
        then vDi:i+15 ← (vA)i:i+15
        else vDi:i+15 ← (vB)i:i+15
end
    
```


Each element of **vmaxuh** is a half word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-49 shows the usage of the **vmaxuh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-49. vmaxuh—Maximum of Eight Unsigned Integer Elements (16-Bit)**


# vmaxuw

# vmaxuw

Vector Maximum Unsigned Word

**vmaxuw**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
  if (vA)i:i+31 ≥ui (vB)i:i+31
    then vDi:i+31 ← (vA)i:i+31
    else vDi:i+31 ← (vB)i:i+31
end
```


Each element of **vmaxuw** is a word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-50 shows the usage of the **vmaxuw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-50. vmaxuw—Maximum of Four Unsigned Integer Elements (32-Bit)**


# vmhaddshs

# vmhaddshs

Vector Multiply High and Add Signed Half Word Saturate

**vmhaddshs**      **vD,vA,vB,vC**

Form: VA



```

do i=0 to 127 by 16
    prod0:31 ← (vA)i:i+15 *si (vB)i:i+15
    temp0:16 ← prod0:16 +int SignExtend((vC)i:i+15,17)
    vDi:i+15 ← SItoSIsat(temp0:16,16)
end
    
```

Each signed-integer half word element in **vA** is multiplied by the corresponding signed-integer half word element in **vB**, producing a 32-bit signed-integer product. Bits 0-16 of the intermediate product are added to the corresponding signed-integer half-word element in **vC** after they have been sign extended to 17-bits. The 16-bit saturated result from each of the eight 17-bit sums is placed in register **vD**.

If the intermediate result is greater than  $(2^{15}-1)$  it saturates to  $(2^{15}-1)$  and if it is less than  $(-2^{15})$  it saturates to  $(-2^{15})$ .

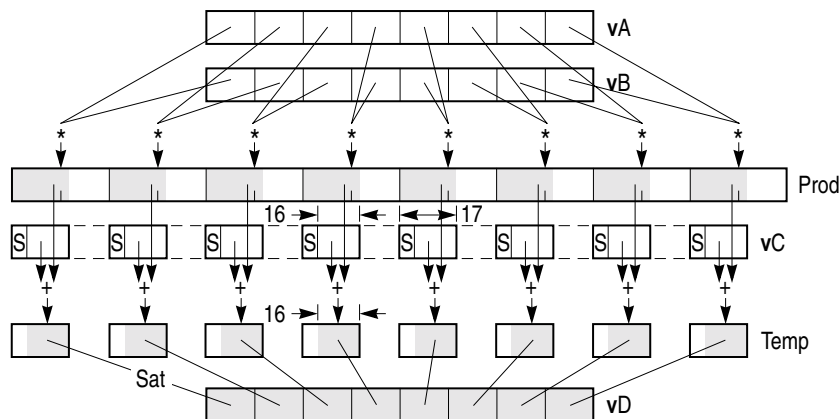
The signed-integer result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT

Figure 6-51 shows the usage of the **vmhaddshs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, **vC**, and **vD**, is 16 bits long.



**Figure 6-51. vmhaddshs—Multiply-High and Add Eight Signed Integer Elements (16-Bit)**


# vmhraddshs

# vmhraddshs

Vector Multiply High Round and Add Signed Half Word Saturate

**vmhraddshs**      **vD,vA,vB,vC**

Form: VA



```

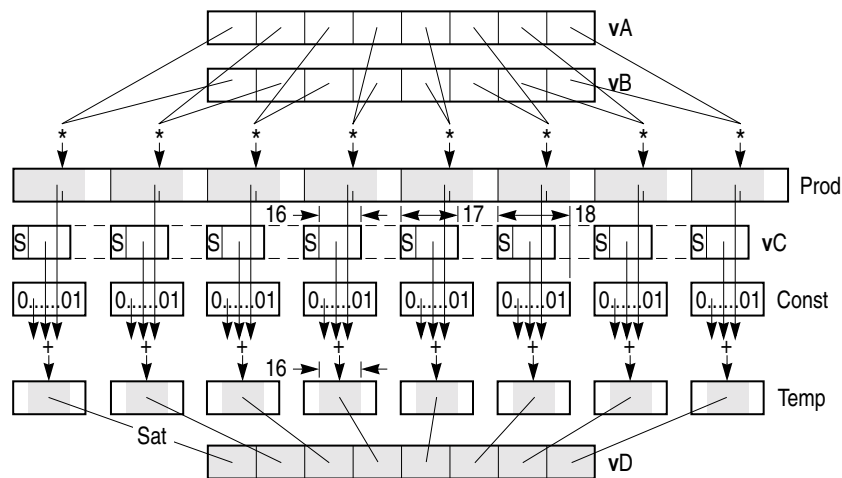
do i=0 to 127 by 16
  prod0:31 ← (vA)i:i+15 *si (vB)i:i+15
  prod0:31 ← prod0:31 +int 0x0000_4000
  temp0:16 ← prod0:16 +int SignExtend((vC)i:i+15,17)
  (vD)i:i+15 ← SItoSIsat(temp0:16,16)
end
  
```

Each signed integer halfword element in register **vA** is multiplied by the corresponding signed integer halfword element in register **vB**, producing a 32-bit signed integer product. The value 0x0000\_4000 is added to the product, producing a 32-bit signed integer sum. Bits 0—16 of the sum are added to the corresponding signed integer halfword element in register **vD**.

If the intermediate result is greater than  $(2^{15}-1)$  it saturates to  $(2^{15}-1)$  and if it is less than  $(-2^{15})$  it saturates to  $(-2^{15})$ .

The signed integer result is and placed into the corresponding halfword element of register **vD**.

Figure 6-52 shows the usage of the **vmhraddshs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, **vC**, and **vD**, is 16 bits long.



**Figure 6-52. vmhraddshs—Multiply-High Round and Add Eight Signed Integer Elements (16-Bit)**


# vminfp

Vector Minimum Floating Point

# vminfp

**vminfp****vD,vA,vB**

Form: VX



```

do i=0 to 127 by 32
  if (vA)i:i+31 <fp (vB)i:i+31
    then vDi:i+31 ← (vA)i:i+31
    else vDi:i+31 ← (vB)i:i+31
end


```

Each single-precision floating-point word element in register **vA** is compared to the corresponding single-precision floating-point word element in register **vB**. The smaller of the two single-precision floating-point values is placed into the corresponding word element of register **vD**.

The minimum of + 0.0 and - 0.0 is - 0.0. The minimum of any value and a NaN is a QNaN.

If VSCR[NJ] = 1, every denormalized operand element is truncated to 0 before the comparison is made.

Figure 6-53 shows the usage of the **vminfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-53. vminfp—Minimum of Four Floating-Point Elements (32-Bit)**


# vminsb

Vector Minimum Signed Byte

# vminsb

**vminsb**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    if (vA)i:i+7 <si (vB)i:i+7
        then vDi:i+7 ← (vA)i:i+7
        else vDi:i+7 ← (vB)i:i+7
end
```


Each element of **vminsb** is a byte.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-54 shows the usage of the **vminsb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-54. vminsb—Minimum of Sixteen Signed Integer Elements (8-Bit)**



# vminsh


Vector Minimum Signed Half Word

# vminsh

**vminsh**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
    if (vA)i:i+15 <si (vB)i:i+15
        then vDi:i+15 ← (vA)i:i+15
        else vDi:i+15 ← (vB)i:i+15
end
```


Each element of **vminsh** is a half word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-55 shows the usage of the **vminsh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-55. vminsh—Minimum of Eight Signed Integer Elements (16-Bit)**

# vminsw


Vector Minimum Signed Word

# vminsw

**vminsw**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    if (vA)i:i+31 <si (vB)i:i+31
        then vDi:i+31 ← (vA)i:i+31
        else vDi:i+31 ← (vB)i:i+31
end
```


Each element of **vminsw** is a word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-56 shows the usage of the **vminsw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-56. vminsw—Minimum of Four Signed Integer Elements (32-Bit)**

# vminub


Vector Minimum Unsigned Byte

# vminub

**vminub**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 8
    if (vA)i:i+7 <ui (vB)i:i+7
        then vDi:i+7 ← (vA)i:i+7
        else vDi:i+7 ← (vB)i:i+7
end
    
```


Each element of **vminub** is a byte.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-57 shows the usage of the **vminub** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-57. vminub—Minimum of Sixteen Unsigned Integer Elements (8-Bit)**

# vminuh


Vector Minimum Unsigned Half Word

# vminuh

**vminuh**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
  if (vA)i:i+15 <ui (vB)i:i+15
    then vDi:i+15 ← (vA)i:i+15
    else vDi:i+15 ← (vB)i:i+15
end
```


Each element of **vminuh** is a half word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-58 shows the usage of the **vminuh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-58. vminuh—Minimum of Eight Unsigned Integer Elements (16-Bit)**

# vminuw


Vector Minimum Unsigned Word

# vminuw

**vminuw**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    if (vA)i:i+31 <ui (vB)i:i+31
        then vDi:i+31 ← (vA)i:i+31
        else vDi:i+31 ← (vB)i:i+31
end
```


Each element of **vminuw** is a word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-59 shows the usage of the **vminuw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-59. vminuw—Minimum of Four Unsigned Integer Elements (32-Bit)**


# vmladduhm

# vmladduhm

Vector Multiply Low and Add Unsigned Half Word Modulo

**vmladduhm**      **vD,vA,vB,vC**

Form: VA



```
do i=0 to 127 by 16
    prod0:31 ← (vA)i:i+15 *ui (vB)i:i+15
    vDi:i+15 ← prod0:31 +int (vC)i:i+15
end
```


Each integer half-word element in **vA** is multiplied by the corresponding integer half-word element in **vB**, producing a 32-bit integer product. The product is added to the corresponding integer half-word element in **vC**. The integer result is placed into the corresponding half-word element of **vD**.

Note that **vmladduhm** can be used for unsigned or signed integers.

Other registers altered:

- None

Figure 6-60 shows the usage of the **vmladduhm** instruction. Each of the eight elements in the vectors, **vA**, **vB**, **vC**, and **vD**, is 16 bits long.



**Figure 6-60. vmladduhm—Multiply-Add of Eight Integer Elements (16-Bit)**


# vmrghb

Vector Merge High Byte

# vmrghb

**vmrghb****vD,vA,vB**

Form: VX



```
do i=0 to 63 by 8
```

```
    vDi*2:(i*2)+15 ← (vA)i:i+7 || (vB)i:i+7
```

```
end
```


Each element of **vmrghb** is a byte.

The elements in the high-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the high-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-61 shows the usage of the **vmrghb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-61. vmrghb—Merge Eight High-Order Elements (8-Bit)**

# vmrghh


Vector Merge High Half word

# vmrghh

**vmrghh**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 16
    vDi*2:(i*2)+31 ← (vA)i:i+15 || (vB)i:i+15
end
```


Each element of **vmrghh** is a half word.

The elements in the high-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the high-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-62 shows the usage of the **vmrghh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-62. vmrghh—Merge Four High-Order Elements (16-Bit)**



# vmrghw


Vector Merge High Word

# vmrghw

**vmrghw**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 32
    vDi*2:(i*2)+63 ← (vA)i:i+31 || (vB)i:i+31
end
```


Each element of **vmrghw** is a word.

The elements in the high-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the high-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-63 shows the usage of the **vmrghw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-63. vmrghw—Merge Four High-Order Elements (32-Bit)**

# vmrglb


Vector Merge Low Byte

# vmrglb

**vmrglb**

**vD,vA,vB**

Form: VX



do i=0 to 63 by 8

$vD_{i*2:(i*2)+15} \leftarrow (vA)_{i+64:i+71} \parallel (vB)_{i+64:i+71}$

end


Each element offer **vmrglb** is a byte.

The elements in the low-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the low-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-64 shows the usage of the **vmrglb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-64. vmrglb—Merge Eight Low-Order Elements (8-Bit)**

# vmrglh


Vector Merge Low Half Word

# vmrglh

**vmrglh**

**vD,vA,vB**

Form: VX



do i=0 to 63 by 16

$$vD_{i*2:(i*2)+31} \leftarrow (vA)_{i+64:i+79} \parallel (vB)_{i+64:i+79}$$

end


Each element of **vmrglh** is a half word.

The elements in the low-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the low-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-65 shows the usage of the **vmrglh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-65. vmrglh—Merge Four Low-Order Elements (16-Bit)**

# vmrglw


Vector Merge Low Word

# vmrglw

**vmrglw**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 32
    vDi*2:(i*2)+63 ← (vA)i+64:i+95 || (vB)i+64:i+95
end
```


Each element of **vmrglw** is a word.

The elements in the low-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the low-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None

Figure 6-66 shows the usage of the **vmrglw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-66. vmrglw—Merge Four Low-Order Elements (32-Bit)**


# vmsummbm

# vmsummbm

Vector Multiply Sum Mixed-Sign Byte Modulo

**vmsummbm**      **vD,vA,vB,vC**

Form: VA



```

do i=0 to 127 by 32
    temp0:31 ← (vC)i:i+31
    do j=0 to 31 by 8
        prod0:15 ← (vA)i+j:i+j+7 *sui (vB)i+j:i+j+7
        temp0:31 ← temp0:31 +int SignExtend(prod0:15, 32)
    end
    vDi:i+31 ← temp0:31
end
    
```


For each word element in **vC** the following operations are performed in the order shown.

- Each of the four signed-integer byte elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**, producing a signed-integer 16-bit product.
- The signed-integer modulo sum of these four products is added to the signed-integer word element in **vC**.
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-67 shows the usage of the **vmsummbm** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, are 8 bits long. Each of the four elements in the vectors, **vC** and **vD** are 32 bits long.



**Figure 6-67. vmsummbm—Multiply-Sum of Integer Elements (8-Bit to 32-Bit)**


# vmsumshm

# vmsumshm

Vector Multiply Sum Signed Half Word Modulo

**vmsumshm**      **vD,vA,vB,vC**

Form: VA



```

do i=0 to 127 by 32
    temp0:31 ← (vC)i:i+31
    do j=0 to 31 by 16
        prod0:31 ← (vA)i+j:i+j+15 *si (vB)i+j:i+j+15
        temp0:31 ← temp0:31 +int prod0:31
        vDi:i+31 ← temp0:31
    end
end
    
```


For each word element in **vC** the following operations are performed in the order shown.

- Each of the two signed-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding signed-integer half-word element in **vB**, producing a signed-integer 32-bit product.
- The signed-integer modulo sum of these two products is added to the signed-integer word element in **vC**.
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-68 shows the usage of the **vmsumshm** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, are 16 bits long. Each of the four elements in the vectors, **vC** and **vD** are 32 bits long.



**Figure 6-68. vmsumshm—Multiply-Sum of Signed Integer Elements (16-Bit to 32-Bit)**

# vmsumshs


# vmsumshs

Vector Multiply Sum Signed Half Word Saturate

vmsumshs

vD,vA,vB,vC

Form: VA



```

do i=0 to 127 by 32
    temp0:33 ← SignExtend((vC)i:i+31,34)
    do j=0 to 31 by 16
        prod0:31 ← (vA)i+j:i+j+15 *si (vB)i+j:i+j+15
        temp0:33 ← temp0:33 +int SignExtend(prod0:31,34)
        vDi:i+31 ← SItoSIsat(temp0:33,32)
    end
end
end
    
```

For each word element in vC the following operations are performed in the order shown.

- Each of the two signed-integer half-word elements in the corresponding word element of vA is multiplied by the corresponding signed-integer half-word element in vB, producing a signed-integer 32-bit product.
- The signed-integer sum of these two products is added to the signed-integer word element in vC.
- If this intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of vD.

Other registers altered:

- SAT

Figure 6-69 shows the usage of the vmsumshs instruction. Each of the eight elements in the vectors, vA, and vB, are 16 bits long. Each of the four elements in the vectors, vC and vD are 32 bits long.




Figure 6-69. vmsumshs—Multiply-Sum of Signed Integer Elements (16-Bit to 32-Bit)


# vmsumubm

# vmsumubm

Vector Multiply Sum Unsigned Byte Modulo

**vmsumubm**      **vD,vA,vB,vC**

Form: VA



```

do i=0 to 127 by 32
    temp0:31 ← (vC)i:i+31
    do j=0 to 31 by 8
        prod0:15 ← (vA)i+j:i+j+7 *ui (vB)i+j:i+j+7
        temp0:32 ← temp0:32 +int ZeroExtend(prod0:15, 32)
        vDi:i+31 ← temp0:31
    end
end
end
    
```


For each word element in **vC** the following operations are performed in the order shown.

- Each of the four unsigned-integer byte elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**, producing an unsigned-integer 16-bit product.
- The unsigned-integer modulo sum of these four products is added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-70 shows the usage of the **vmsumubm** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, are 8 bits long. Each of the four elements in the vectors, **vC** and **vD** are 32 bits long.



**Figure 6-70. vmsumubm—Multiply-Sum of Unsigned Integer Elements (8-Bit to 32-Bit)**




**vmsumuhm**

Vector Multiply Sum Unsigned Half Word Modulo

**vmsumuhm****vmsumuhm**      **vD,vA,vB,vC**

Form: VA



```

do i=0 to 127 by 32
  temp0:31 ← (vC)i:i+31
  do j=0 to 31 by 16
    prod0:31 ← (vA)i+j:i+j+15 *ui (vB)i+j:i+j+15
    temp0:31 ← temp0:31 +int prod0:31
    vDi:i+31 ← temp2:33
  end
end
end

```


For each word element in **vC** the following operations are performed in the order shown.

- Each of the two unsigned-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**, producing a unsigned-integer 32-bit product.
- The unsigned-integer sum of these two products is added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-71 shows the usage of the **vmsumuhm** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, are 16 bits long. Each of the four elements in the vectors, **vC** and **vD** are 32 bits long.



**Figure 6-71. vmsumuhm—Multiply-Sum of Unsigned Integer Elements (16-Bit to 32-Bit)**

# vmsumuhs


# vmsumuhs

Vector Multiply Sum Unsigned Half Word Saturate

**vmsumuhs**

**vD,vA,vB,vC**

Form: VA



```
do i=0 to 127 by 32
    temp0:33 ← ZeroExtend((vC)i:i+31,34)
    do j=0 to 31 by 16
        prod0:31 ← (vA)i+j:i+j+15 *ui (vB)i+j:i+j+15
        temp0:33 ← temp0:33 +int ZeroExtend(prod0:31,34)
        vDi:i+31 ← UItoUISat(temp0:33,32)
    end
end
```


For each word element in **vC** the following operations are performed in the order shown.

- Each of the two unsigned-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**, producing an unsigned-integer 32-bit product.
- The unsigned-integer sum of these two products is saturate-added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT

Figure 6-72 shows the usage of the **vmsumuhs** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, are 16 bits long. Each of the four elements in the vectors, **vC** and **vD** are 32 bits long.



**Figure 6-72. vmsumuhs—Multiply-Sum of Unsigned Integer Elements (16-Bit to 32-Bit)**


# vmulesb

# vmulesb

Vector Multiply Even Signed Byte

**vmulesb**                      **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 16
    prod0:15 ← (vA)i:i+7 *si (vB)i:i+7
    vDi:i+15 ← prod0:15
end
```

Each even-numbered signed-integer byte element in **vA** is multiplied by the corresponding signed-integer byte element in **vB**. The eight 16-bit signed-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None

Figure 6-73 shows the usage of the **vmulesb** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, is 8 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-73. vmulesb—Even Multiply of Eight Signed Integer Elements (8-Bit)**

# vmulesh


# vmulesh

Vector Multiply Even Signed Half Word

**vmulesh**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
    prod0:31 ← (vA)i:i+15 *si (vB)i:i+15
    vDi:i+31 ← prod0:31
end
```

Each even-numbered signed-integer half-word element in **vA** is multiplied by the corresponding signed-integer half-word element in **vB**. The four 32-bit signed-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None

Figure 6-74 shows the usage of the **vmulesh** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the four elements in the vector **vD**, is 32 bits long.



**Figure 6-74. vmulesb—Even Multiply of Four Signed Integer Elements (16-Bit)**


# vmuleub

# vmuleub

Vector Multiply Even Unsigned Byte

**vmuleub**                      **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 16
    prod0:15 ← (vA)i:i+7 *ui (vB)i:i+7
    (vD)i:i+15 ← prod0:15
end
```

Each even-numbered unsigned-integer byte element in register **vA** is multiplied by the corresponding unsigned-integer byte element in register **vB**. The eight 16-bit unsigned-integer products are placed, in the same order, into the eight halfwords of register **vD**.

Other registers altered:

- None

Figure 6-75 shows the usage of the **vmuleub** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, is 8 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.




**Figure 6-75. vmuleub—Even Multiply of Eight Unsigned Integer Elements (8-Bit)**

# vmuleuh

# vmuleuh

Vector Multiply Even Unsigned Half Word

**vmuleuh**                      **vD,vA,vB**                      Form: VX




```
do i=0 to 127 by 32
    prod0:31 ← (vA)i:i+15 *ui (vB)i:i+15
    (vD)i:i+31 ← prod0:31
end
```

Each even-numbered unsigned-integer halfword element in register **vA** is multiplied by the corresponding unsigned-integer halfword element in register **vB**. The four 32-bit unsigned-integer products are placed, in the same order, into the four words of register **vD**.

Other registers altered:

- None

Figure 6-76 shows the usage of the **vmuleuh** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the four elements in the vector **vD**, is 32 bits long.



**Figure 6-76. vmuleuh—Even Multiply of Four Unsigned Integer Elements (16-Bit)**


# vmulosb

# vmulosb

Vector Multiply Odd Signed Byte

**vmulosb**                      **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 16
    prod0:15 ← (vA)i+8:i+15 *si (vB)i+8:i+15
    vDi:i+15 ← prod0:15
end
```

Each odd-numbered signed-integer byte element in **vA** is multiplied by the corresponding signed-integer byte element in **vB**. The eight 16-bit signed-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None

Figure 6-77 shows the usage of the **vmulosb** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, is 8 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-77. vmulosb—Odd Multiply of Eight Signed Integer Elements (8-Bit)**


# vmulosh

# vmulosh

Vector Multiply Odd Signed Half Word

**vmulosh**                      **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
    prod0:31 ← (vA)i+16:i+31 *si (vB)i+16:i+31
    vDi:i+31 ← prod0:31
end
```

Each odd-numbered signed-integer half-word element in **vA** is multiplied by the corresponding signed-integer half-word element in **vB**. The four 32-bit signed-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None

Figure 6-78 shows the usage of the **vmuleuh** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the four elements in the vector **vD**, is 32 bits long.



**Figure 6-78. vmuleuh—Odd Multiply of Four Unsigned Integer Elements (16-Bit)**




# vmuloub

# vmuloub

Vector Multiply Odd Unsigned Byte

**vmuloub**                      **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 8
    prod0:15 ← (vA)i+8:i+15 *ui (vB)i+n:i+15
    vDi:i+15 ← prod0:15
end
```

Each odd-numbered unsigned-integer byte element in **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**. The eight 16-bit unsigned-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None

Figure 6-79 shows the usage of the **vmuloub** instruction. Each of the sixteen elements in the vectors, **vA**, and **vB**, is 8 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.




**Figure 6-79. vmuloub—Odd Multiply of Eight Unsigned Integer Elements (8-Bit)**

# vmulouh

# vmulouh

Vector Multiply Odd Unsigned Half Word

**vmulouh**                      **vD,vA,vB**                      Form: VX




```
do i=0 to 127 by 16
    prod0:31 ← (vA)i+16:i+31 *ui (vB)i+n:i+31
    vDi:i+31 ← prod0:31
end
```

Each odd-numbered unsigned-integer half-word element in **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**. The four 32-bit unsigned-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None

Figure 6-80 shows the usage of the **vmulouh** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the four elements in the vector **vD**, is 32 bits long.



**Figure 6-80. vmulouh—Odd Multiply of Four Unsigned Integer Elements (16-Bit)**

# vnmsubfp


# vnmsubfp

Vector Negative Multiply-Subtract Floating Point

**vnmsubfp**

**vD,vA,vC,vB**

Form: VA



do i=0 to 127 by 32

$$vD_{i:i+31} \leftarrow -\text{RndToNearFP32}((vA)_{i:i+31} *_{fp} (vC)_{i:i+31}) -_{fp} (vB)_{i:i+31}$$

end

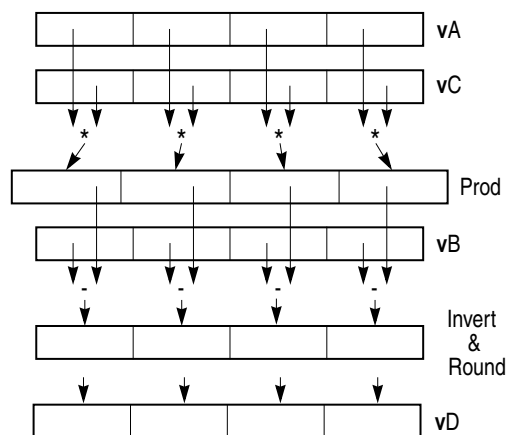
Each single-precision floating-point word element in **vA** is multiplied by the corresponding single-precision floating-point word element in **vC**. The corresponding single-precision floating-point word element in **vB** is subtracted from the product. The sign of the difference is inverted. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

Note that only one rounding occurs in this operation. Also note that a QNaN result is not negated.

Other registers altered:

- None

Figure 6-81 shows the usage of the **vnmsubfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-81. vnmsubfp—Negative Multiply-Subtract of Four Floating-Point Elements (32-Bit)**

# vnor


Vector Logical NOR

# vnor

**vnor**

**vD,vA,vB**

Form: VX



$$vD \leftarrow \neg((vA) \mid (vB))$$

The contents of **vA** are bitwise ORed with the contents of **vB** and the complemented result is placed into **vD**.


Other registers altered:

- None

Simplified mnemonics:

**vnot vD, vS** equivalent to **vnor vD, vS, vS**

Figure 6-82 shows the usage of the **vnor** instruction.



**Figure 6-82. vnor—Bitwise NOR of 128-bit Vector**

# vor


Vector Logical OR

# VOR

**vor**

**vD,vA,vB**

Form: VX



$$vD \leftarrow (vA) \mid (vB)$$

The contents of vA are ORed with the contents of vB and the result is placed into vD.

Other registers altered:


- None

Simplified mnemonics:

**vmr vD, vS**

equivalent to **vor vD, vS, vS**

Figure 6-83 shows the usage of the vor instruction.




**Figure 6-83. vor—Bitwise OR of 128-bit Vector**

# vperm

Vector Permute

# vperm

**vperm**                      **vD,vA,vB,vC**                      Form: VA



```
temp0:255 ← (vA) || (vB)
do i=0 to 127 by 8
    b ← (vC)i+3:i+7 || 0b000
    vDi:i+7 ← tempb:b+7
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**. For each integer *i* in the range 0–15, the contents of the byte element in the source vector specified in bits 3–7 of byte element *i* in **vC** are placed into byte element *i* of **vD**.

Other registers altered:

- None

Programming note: See the programming notes with the Load Vector for Shift Left and Load Vector for Shift Right instructions for examples of usage on the **vperm** instruction.

Figure 6-84 shows the usage of the **vperm** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, **vC**, and **vD**, is 8 bits long.



**Figure 6-84. vperm—Concatenate Sixteen Integer Elements (8-Bit)**

# vpxpx


Vector Pack Pixel32

# vpxpx

vpxpx

vD,vA,vB

Form: VX




```

do i=0 to 63 by 16
    vDi ← (vA)i*2+7
    vDi+1:i+5 ← (vA)(i*2)+8:(i*2)+12
    vDi+6:i+10 ← (vA)(i*2)+16:(i*2)+20
    vDi+11:i+15 ← (vA)(i*2)+24:(i*2)+28
    vDi+64 ← (vB)(i*2)+7
    vDi+65:i+69 ← (vB)(i*2)+8:(i*2)+12
    vDi+70:i+74 ← (vB)(i*2)+16:(i*2)+20
    vDi+75:i+79 ← (vB)(i*2)+24:(i*2)+28
end
    
```

The source vector is the concatenation of the contents of vA followed by the contents of vB. Each 32-bit word element in the source vector is packed to produce a 16-bit half-word value as described below and placed into the corresponding half-word element of vD. A word is packed to 16 bits by concatenating, in order, the following bits.

- bit 7 of the first byte (bit 7 of the word)
- bits 0–4 of the second byte (bits 8–12 of the word)
- bits 0–4 of the third byte (bits 16–20 of the word)
- bits 0–4 of the fourth byte (bits 24–28 of the word)

Figure 6-85 shows which bits of the source word are packed to form the half word in the destination register.




**Figure 6-85. How a Word is Packed to a Half Word**

Other registers altered:

- None

Programming note: Each source word can be considered to be a 32-bit pixel consisting of four 8-bit channels. Each target half-word can be considered to be a 16-bit pixel consisting of one 1-bit channel and three 5-bit channels. A channel can be used to specify the intensity of a particular color, such as red, green, or blue, or to provide other information needed by the application.

Figure 6-86 shows the usage of the **vpkpx** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-86. vpkpx—Pack Eight Elements (32-Bit) to Eight Elements (16-Bit)**



# vpkshss


Vector Pack Signed Half Word Signed Saturate

# vpkshss

vpkshss

vD,vA,vB

Form: VX



```
do i=0 to 63 by 8
```

```
  vDi:i+7 ← SItoSIsat((vA)i*2:(i*2)+15, 8)
```

```
  vDi+64:i+71 ← SItoSIsat((vB)i*2:(i*2)+15, 8)
```

```
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer half-word element in the source vector is converted to an 8-bit signed integer. If the value of the element is greater than  $(2^7 - 1)$  the result saturates to  $(2^7 - 1)$  and if the value is less than  $-2^7$  the result saturates to  $-2^7$ . The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT

Figure 6-87 shows the usage of the **vpkshss** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the sixteen elements in the vector **vD**, is 8 bits long.



**Figure 6-87. vpkshss—Pack Sixteen Signed Integer Elements (16-Bit) to Sixteen Signed Integer Elements (8-Bit)**


# vpkshus

Vector Pack Signed Half Word Unsigned Saturate

# vpkshus

**vpkshus**                      **vD,vA,vB**

Form: VX



```
do i=0 to 63 by 8
    vDi:i+7 ← SItOUIsat((vA)i*2:(i*2)+7, 8)
    vDi+64:i+71 ← SItOUIsat((vB)i*2:(i*2)+7, 8)
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer half-word element in the source vector is converted to an 8-bit unsigned integer. If the value of the element is greater than  $(2^8 - 1)$  the result saturates to  $(2^8 - 1)$  and if the value is less than 0 the result saturates to 0. The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT

Figure 6-88 shows the usage of the **vpkshus** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the sixteen elements in the vector **vD**, is 8 bits long.



**Figure 6-88. vpkshus—Pack Sixteen Signed Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit)**

# vpkswss


Vector Pack Signed Word Signed Saturate

# vpkswss

vpkswss

vD,vA,vB

Form: VX



```
do i=0 to 63 by 16
```

```
  vDi:i+15 ← SItoSIsat((vA)i*2:(i*2)+31, 16)
  vDi+64:i+79 ← SItoSIsat((vB)i*2:(i*2)+31, 16)
```

```
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer word element in the source vector is converted to a 16-bit signed integer half word. If the value of the element is greater than  $(2^{15} - 1)$  the result saturates to  $(2^{15} - 1)$  and if the value is less than  $-2^{15}$  the result saturates to  $-2^{15}$ . The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT

Figure 6-89 shows the usage of the **vpkswss** instruction. Each of the four elements in the vectors, **vA**, and **vB**, is 32 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-89. vpkswss—Pack Eight Signed Integer Elements (32-Bit) to Eight Signed Integer Elements (16-Bit)**

# vpkswus


Vector Pack Signed Word Unsigned Saturate

# vpkswus

**vpkswus**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 16
    vDi:i+15 ← SItOUIsat((vA)i*2:(i*2)+31, 16)
    vDi+64:i+79 ← SItOUIsat((vB)i*2:(i*2)+31, 16)
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer word element in the source vector is converted to a 16-bit unsigned integer. If the value of the element is greater than  $(2^{16} - 1)$  the result saturates to  $(2^{16} - 1)$  and if the value is less than 0 the result saturates to 0. The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT

Figure 6-90 shows the usage of the **vpkswus** instruction. Each of the four elements in the vectors, **vA**, and **vB**, is 32 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-90. vpkswus—Pack Eight Signed Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit)**

# vpkuhum


Vector Pack Unsigned Half Word Unsigned Modulo

# vpkuhum

**vpkuhum**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 8
    vDi:i+7 ← (vA)(i*2)+8:(i*2)+15
    vDi+64:i+71 ← (vB)(i*2)+8:(i*2)+15
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

The low-order byte of each half-word element in the source vector is placed into the corresponding byte element of **vD**.

Other registers altered:

- None

Figure 6-91 shows the usage of the **vpkuhum** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the sixteen elements in the vector **vD**, is 8 bits long.



**Figure 6-91. vpkuhum—Pack Sixteen Unsigned Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit)**

# vpkuhus


Vector Pack Unsigned Half Word Unsigned Saturate

# vpkuhus

**vpkuhus**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 8
    vDi:i+7 ← UItoUISat((vA)i*2:(i*2)+15, 8)
    vDi+64:i+71 ← UItoUISat((vB)i*2:(i*2)+15, 8)
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each unsigned integer half-word element in the source vector is converted to an 8-bit unsigned integer. If the value of the element is greater than  $(2^8 - 1)$  the result saturates to  $(2^8 - 1)$ . The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT

Figure 6-92 shows the usage of the **vpkuhus** instruction. Each of the eight elements in the vectors, **vA**, and **vB**, is 16 bits long. Each of the sixteen elements in the vector **vD**, is 8 bits long.



**Figure 6-92. vpkuhus—Pack Sixteen Unsigned Integer Elements (16-Bit) to Sixteen Unsigned Integer Elements (8-Bit)**

# vpkuwum


Vector Pack Unsigned Word Unsigned Modulo

# vpkuwum

**vpkuwum**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 16
    vDi:i+15 ← (vA)(i*2)+16:(i*2)+31
    vDi+64:i+79 ← (vB)(i*2)+16:(i*2)+31
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

The low-order half-word of each word element in the source vector is placed into the corresponding half-word element of **vD**.

Other registers altered:

- None

Figure 6-93 shows the usage of the **vpkuwum** instruction. Each of the four elements in the vectors, **vA**, and **vB**, is 32 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-93. vpkuwum—Pack Eight Unsigned Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit)**

# vpkuwus


# vpkuwus

Vector Pack Unsigned Word Unsigned Saturate

**vpkuwus**

**vD,vA,vB**

Form: VX



```
do i=0 to 63 by 16
    vDi:i+15 ← UItoUISat((vA)i*2:(i*2)+31, 16)
    vDi+64:i+79 ← UItoUISat((vB)i*2:(i*2)+31, 16)
end
```


Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each unsigned integer word element in the source vector is converted to a 16-bit unsigned integer. If the value of the element is greater than  $(2^{16} - 1)$  the result saturates to  $(2^{16} - 1)$ . The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT

Figure 6-94 shows the usage of the **vpkuwus** instruction. Each of the four elements in the vectors, **vA**, and **vB**, is 32 bits long. Each of the eight elements in the vector **vD**, is 16 bits long.



**Figure 6-94. vpkuwum—Pack Eight Unsigned Integer Elements (32-Bit) to Eight Unsigned Integer Elements (16-Bit)**




**vrefp**

Vector Reciprocal Estimate Floating Point

**vrefp****vrefp****vD,vB**

Form: VX



```

do i=0 to 127 by 32
  x ← (vB)i:i+31
  vDi:i+31 ← 1/x
end

```

The single-precision floating-point estimate of the reciprocal of each single-precision floating-point element in **vB** is placed into the corresponding element of **vD**.

For results that are not a +0, -0, +∞, -∞, or QNaN, the estimate has a relative error in precision no greater than one part in 4096, that is:

$$\left| \frac{\text{estimate} - 1/x}{1/x} \right| \leq \frac{1}{4096}$$

where  $x$  is the value of the element in **vB**. Note that the value placed into the element of **vD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the element in **vB** is summarized below in Table 6-7.

**Table 6-7. Special Values of the Element in vB**


Value	Result
-∞	-0
-0	-∞
+0	+∞
+∞	+0
NaN	QNaN

If  $VSCR[NJ] = 1$ , every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign.

Other registers altered:

- None

Figure 6-95 shows the usage of the **vrefp** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-95. vrefp—Reciprocal Estimate of Four Floating-Point Elements (32-Bit)**

# vrfim


# vrfim

Vector Round to Floating-Point Integer toward Minus Infinity

**vrfim**

**vD,vB**

Form: VX




```
do i=0 to 127 by 32
    vDi:i+31 ← RndToFPInt32Floor((vB)i:i+31)
end
```

Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer, using the rounding mode Round toward -Infinity, and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-96 shows the usage of the **vrfim** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-96. vrfim— Round to Minus Infinity of Four Floating-Point Integer Elements (32-Bit)**

# vrfn


# vrfn

Vector Round to Floating-Point Integer Nearest

**vrfn**

**vD,vB**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← RndToFPInt32Near((vB)i:i+31)
end
```


Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer, using the rounding mode Round to Nearest, and placed into the corresponding word element of **vD**.

Note the result is independent of VSCR[NJ].

Other registers altered:

- None

Figure 6-97 shows the usage of the **vrfn** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-97. vrfn—Nearest Round to Nearest of Four Floating-Point Integer Elements (32-Bit)**

# vrfip


Vector Round to Floating-Point Integer toward Plus Infinity

# vrfip

vrfip

vD,vB

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← RndToFPInt32Ceil((vB)i:i+31)
end
```


Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer, using the rounding mode Round toward +Infinity, and placed into the corresponding word element of **vD**.

If **VSCR[NJ] = 1**, every denormalized operand element is truncated to 0 before the comparison is made.

Other registers altered:

- None

Figure 6-98 shows the usage of the **vrfip** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-98. vrfip—Round to Plus Infinity of Four Floating-Point Integer Elements (32-Bit)**

# vrfiz


# vrfiz

Vector Round to Floating-Point Integer toward Zero

**vrfiz**

**vD,vB**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← RndToFPInt32Trunc((vB)i:i+31)
end
```


Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer, using the rounding mode Round toward Zero, and placed into the corresponding word element of **vD**.

Note, the result is independent of VSCR[NJ].

Other registers altered:

- None

Figure 6-99 shows the usage of the **vrfiz** instruction. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-99. vrfiz—Round-to-Zero of Four Floating-Point Integer Elements (32-Bit)**

# vrlb


Vector Rotate Left Integer Byte

# vrlb

**vrlb**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 8
    sh ← (vB)i+5:i+7
    vDi:i+7 ← ROTL((vA)i:i+7, sh)
end
```

Each element is a byte. Each element in **vA** is rotated left by the number of bits specified in the low-order 3 bits of the corresponding element in **vB**. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-100 shows the usage of the **vrlb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-100. vrlb—Left Rotate of Sixteen Integer Elements (8-Bit)**

# vrlh


Vector Rotate Left Integer Half Word

# vrlh

**vrlh**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
  sh ← (vB)i+12:i+15
  vDi:i+15 ← ROTL((vA)i:i+15,sh)
end
```


Each element is a half word

Each element in **vA** is rotated left by the number of bits specified in the low-order 4 bits of the corresponding element in **vB**. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-101 shows the usage of the **vrlh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-101. vrlh—Left Rotate of Eight Integer Elements (16-Bit)**



# vrlw


Vector Rotate Left Integer Word

# vrlw

**vrlw**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
    sh ← (vB)i+27:i+31
    vDi:i+31 ← ROTL((vA)i:i+31,sh)
end
```

Each element is a word. Each element in **vA** is rotated left by the number of bits specified in the low-order 5 bits of the corresponding element in **vB**. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-102 shows the usage of the **vrlw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-102. vrlw—Left Rotate of Four Integer Elements (32-Bit)**

# vrsqrtefp


# vrsqrtefp

Vector Reciprocal Square Root Estimate Floating Point

**vrsqrtefp**

**vD,vB**

Form: VX



```
do i=0 to 127 by 32
    x ← (vB)i:i+31
    vDi:i+31 ← 1 ÷fp (√fp(x))
end
```

The single-precision estimate of the reciprocal of the square root of each single-precision element in **vB** is placed into the corresponding word element of **vD**. The estimate has a relative error in precision no greater than one part in 4096, as explained below:

$$\left| \frac{\text{estimate} - 1/\sqrt{x}}{1/\sqrt{x}} \right| \leq \frac{1}{4096}$$

where  $x$  is the value of the element in **vB**. Note that the value placed into the element of **vD** may vary between implementations and between different executions on the same implementation. Operation with various special values of the element in **vB** is summarized below in Table 6-8.


**Table 6-8. Special Values of the Element in vB**

Value	Result	Value	Result
$-\infty$	QNaN	+0	$+\infty$
less than 0	QNaN	$+\infty$	+0
-0	$-\infty$	NaN	QNaN

Other registers altered:

- None

Figure 6-103 shows the usage of the **vrsqrtefp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-103. vrsqrtefp—Reciprocal Square Root Estimate of Four Floating-Point Elements (32-Bit)**


# vsel

Vector Conditional Select

# vsel

**vsel**                      **vD,vA,vB,vC**

Form: VA




```
do i=0 to 127
    if (vC)i=0 then vDi ← (vA)i
    else vDi ← (vB)i
end
```

For each bit in **vC** that contains the value 0, the corresponding bit in **vA** is placed into the corresponding bit of **vD**. For each bit in **vC** that contains the value 1, the corresponding bit in **vB** is placed into the corresponding bit of **vD**.

Other registers altered:

- None

Figure 6-104 shows the usage of the **vsel** instruction. Each of the vectors, **vA**, **vB**, **vC**, and **vD**, is 128 bits long.



**Figure 6-104. vsel—Bitwise Conditional Select of Vector Contents(128-bit)**

# vsl


Vector Shift Left

# vsl

vsl

vD,vA,vB

Form: VX



```

sh ← (vB)125:127
t ← 1
do i = 0 to 127 by 8
    t ← t & ((vB)i+5:i+7 = sh)
    if t = 1 then vD ← (vA) <<ui sh
    else vD ← undefined
end
    
```


The contents of vA are shifted left by the number of bits specified in vB[125–127]. Bits shifted out of bit 0 are lost. Zeros are supplied to the vacated bits on the right. The result is placed into vD.

The contents of the low-order three bits of all byte elements in vB must be identical to vB[125–127]; otherwise the value placed into vD is undefined.

Other registers altered:

- None

Figure 6-105 shows the usage of the vsl instruction.



**Figure 6-105. vsl—Shift Bits Left in Vector (128-Bit)**

# vslb


Vector Shift Left Integer Byte

# vslb

**vslb**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 8
    sh ← (vB)i+5:i+7
    vDi:i+7 ← (vA)i:i+7 <<ui sh
end
```

Each element is a byte. Each element in **vA** is shifted left by the number of bits specified in the low-order 3 bits of the corresponding element in **vB**. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-106 shows the usage of the **vslb** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-106. vslb—Shift Bits Left in Sixteen Integer Elements (8-Bit)**


# vsldoi

Vector Shift Left Double by Octet Immediate

# vsldoi

**vsldoi**            **vD, vA, vB, SHB**

Form: VA




$$vD \leftarrow ((vA) \parallel (vB)) \ll_{ui} (SHB \parallel 0b000)$$

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**. Bytes **SHB:SHB+15** of the source vector are placed into **vD**.

Other registers altered:

- None

Figure 6-107 shows the usage of the **vsldoi** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-107. vsldoi—Shift Left by Bytes Specified**

# vslh


Vector Shift Left Integer Half Word

# vslh

**vslh**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 16
    sh ← (vB)i+12:i+15
    vDi:i+15 ← (vA)i:i+15 <<ui sh
end
```

Each element is a half word. Each element in **vA** is shifted left by the number of bits specified in the low-order 4 bits of the corresponding element in **vB**. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-108 shows the usage of the **vslh** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-108. vslh—Shift Bits Left in Eight Integer Elements (16-Bit)**

# vslo


Vector Shift Left by Octet

# vslo

**vslo**

**vD,vA,vB**

Form: VX




```
shb ← (vB)121:124
vD ← (vA) <<ui (shb || 0b000)
```

The contents of **vA** are shifted left by the number of bytes specified in **vB**[121–124]. Bytes shifted out of byte 0 are lost. Zeros are supplied to the vacated bytes on the right. The result is placed into **vD**.

Other registers altered:

- None

Figure 6-109 shows the usage of the **vslo** instruction.



**Figure 6-109. vslo—Left Byte Shift of Vector (128-Bit)**



# vslw


Vector Shift Left Integer Word

# vslw

vslw

vD,vA,vB

Form: VX



```
do i=0 to 127 by 32
    sh ← (vB)i+27:i+31
    vDi:i+31 ← (vA)i:i+31 <<ui sh
end
```

Each element is a word. Each element in vA is shifted left by the number of bits specified in the low-order 5 bits of the corresponding element in vB. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of vD.

Other registers altered:

- None

Figure 6-110 shows the usage of the vslw instruction. Each of the four elements in the vectors, vA, vB, and vD, is 32 bits long.

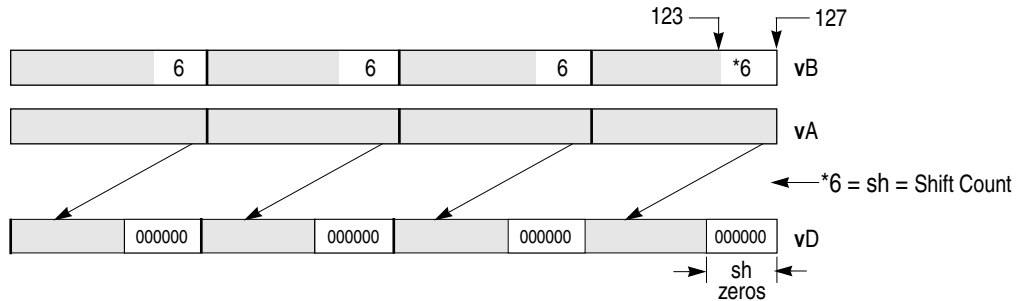


Figure 6-110. vslw—Shift Bits Left in Four Integer Elements (32-Bit)


# vspltb

Vector Splat Byte

# vspltb

**vspltb**                      **vD,vB,UIMM**

Form: VX



```

b ← UIMM*8
do i=0 to 127 by 8
    vDi:i+7 ← (vB)b:b+7
end
    
```

Each element of **vspltb** is a byte.


The contents of element UIMM in **vB** are replicated into each element of **vD**.

Other registers altered:

- None

Programming note: The vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a vector register by a constant).

Figure 6-111 shows the usage of the **vspltb** instruction. Each of the sixteen elements in the vectors **vB** and **vD** is 8 bits long.



**Figure 6-111. vspltb—Copy Contents to Sixteen Elements (8-Bit)**


# vsplth

Vector Splat Half Word

# vsplth

**vsplth****vD,vB,UIMM**

Form: VX



```

b ← UIMM*16
do i=0 to 127 by 16
    vDi:i+15 ← (vB)b:b+15
end

```

Each element of **vsplth** is a half word.


The contents of element UIMM in **vB** are replicated into each element of **vD**.

Other registers altered:

- None

Programming note: The vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a vector register by a constant).

Figure 6-112 shows the usage of the **vsplth** instruction. Each of the eight elements in the vectors **vB** and **vD** is 16 bits long.



**Figure 6-112. vsplth—Copy Contents to Eight Elements (16-Bit)**

# vspltisb


Vector Splat Immediate Signed Byte

# vspltisb

**vspltisb**

vD,SIMM

Form: VX



```
do i=0 to 127 by 8
    vDi:i+7 ← SignExtend(SIMM,8)
end
```


Each element of **vspltisb** is a byte.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of vD.

Other registers altered:

- None

Figure 6-113 shows the usage of the **vspltisb** instruction. Each of the sixteen elements in the vector, vD, is 8 bits long.



**Figure 6-113. vspltisb—Copy Value into Sixteen Signed Integer Elements (8-Bit)**

# vspltish


Vector Splat Immediate Signed Half Word

# vspltish

**vspltish**

**vD,SIMM**

Form: VX



```
do i=0 to 127 by 16
    vDi:i+15 ← SignExtend(SIMM,16)
end
```


Each element of **vspltish** is a half word.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of **vD**.

Other registers altered:

- None

Figure 6-114 shows the usage of the **vspltish** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-114. vspltish—Copy Value to Eight Signed Integer Elements (16-Bit)**

# vspltisw


Vector Splat Immediate Signed Word

# vspltisw

**vspltisw**

**vD,SIMM**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← SignExtend(SIMM,32)
end
```


Each element of **vspltisw** is a word.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of **vD**.

Other registers altered:

- None

Figure 6-115 shows the usage of the **vspltisw** instruction. Each of the four elements in the vector, and **vD**, is 32 bits long.



**Figure 6-115. vspltisw—Copy Value to Four Signed Elements (32-Bit)**

# vspltw


Vector Splat Word

# vspltw

**vspltw**

**vD,vB,UIMM**

Form: VX



```

b ← UIMM*32
do i=0 to 127 by 32
    vDi:i+31 ← (vB)b:b+31
end
    
```

Each element of **vspltw** is a word.


The contents of element UIMM in **vB** are replicated into each element of **vD**.

Other registers altered:

- None

Programming note: The Vector Splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a Vector Register by a constant).

Figure 6-116 shows the usage of the **vspltw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-116. vspltw—Copy contents to Four Elements (32-Bit)**

# VSR


Vector Shift Right

# VSR

**vsr**

**vD,vA,vB**

Form: VX



```

sh ← (vB)125:127
t ← 1
do i = 0 to 127 by 8
    t ← t & ((vB)i+5:i+7 = sh)
    if t = 1 then vD ← (vA) >>ui sh
    else vD ← undefined
end
    
```

Let  $sh = vB[125-127]$ ;  $sh$  is the shift count in bits ( $0 \leq sh \leq 7$ ). The contents of  $vA$  are shifted right by  $sh$  bits. Bits shifted out of bit 127 are lost. Zeros are supplied to the vacated bits on the left. The result is placed into  $vD$ .

The contents of the low-order three bits of all byte elements in register  $vB$  must be identical to  $vB[125-127]$ ; otherwise the value placed into register  $vD$  is undefined.

Other registers altered:

- None

Programming notes:

A pair of **vslo** and **vsl** or **vsro** and **vsr** instructions, specifying the same shift count register, can be used to shift the contents of a vector register left or right by the number of bits (0–127) specified in the shift count register. The following example shifts the contents of  $vX$  left by the number of bits specified in  $vY$  and places the result into  $vZ$ .

```

vslo    VZ,VX,VY
vsl     VZ,VZ,VY
    
```

A double-register shift by a dynamically specified number of bits (0–127) can be performed in six instructions. The following example shifts  $(vW) \parallel (vX)$  left by the number of bits specified in  $vY$  and places the high-order 128 bits of the result into  $vZ$ .

```

vslo    t1,VW,VY #shift high-order reg left
vsl     t1,t1,VY
vsububm t3,V0,VY #adjust shift count ((V0)=0)
vsro    t2,VX,t3 #shift low-order reg right
vsr     t2,t2,t3
vor     VZ,t1,t2 #merge to get final result
    
```



Figure 6-117 shows the usage of the `vsr` instruction. Each of the sixteen elements in the vectors, `vA`, `vB`, and `vD`, is 8 bits long.




Figure 6-117. `vsr`—Shift Bits Right for Vectors (128-Bit)


# vsrab

# vsrab

Vector Shift Right Algebraic Byte

**vsrab**                                  **vD,vA,vB**

Form: VX




```
do i=0 to 127 by 8
    sh ← (vB)i+2:i+7
    vDi:i+7 ← (vA)i:i+7 >>si sh
end
```

Each element is a byte. Each element in **vA** is shifted right by the number of bits specified in the low-order 3 bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-118 shows the usage of the **vsrab** instruction. Each of the sixteen elements in the vectors, **vA**, and **vD**, is 8 bits long.



**Figure 6-118. vsrab—Shift Bits Right in Sixteen Integer Elements (8-Bit)**

# vsrah


# vsrah

Vector Shift Right Algebraic Half Word

**vsrah**

**vD,vA,vB**

Form: VX



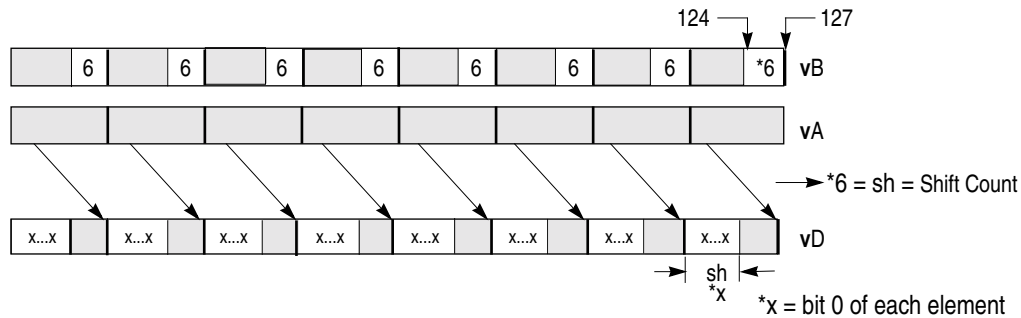
```
do i=0 to 127 by 16
    sh ← (vB)i+12:i+15
    vDi:i+15 ← (vA)i:i+15 >>si sh
end
```

Each element is a half word. Each element in **vA** is shifted right by the number of bits specified in the low-order 4 bits of the corresponding element in **vB**. Bits shifted out of bit 15 of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-119 shows the usage of the **vsrah** instruction. Each of the eight elements in the vectors, **vA**, and **vD**, is 16 bits long.



**Figure 6-119. vsrah—Shift Bits Right for Eight Integer Elements (16-Bit)**

# vsraw


Vector Shift Right Algebraic Word

# vsraw

**vsraw**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
  sh ← (vB)i+27:i+31
  vDi:i+31 ← (vA)i:i+31 >>si sh
end
```

Each element is a word. Each element in **vA** is shifted right by the number of bits specified in the low-order 5 bits of the corresponding element in **vB**. Bits shifted out of bit 31 of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-120 shows the usage of the **vsraw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-120. vsraw—Shift Bits Right in Four Integer Elements (32-Bit)**

# vsrb


Vector Shift Right Byte

# vsrb

**vsrb**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 8
    sh ← (vB)i+5:i+7
    vDi:i+7 ← (vA)i:i+7 >>ui sh
end
```

Each element is a byte. Each element in **vA** is shifted right by the number of bits specified in the low-order 3 bits of the corresponding element in **vB**. Bits shifted out of bit 7 of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-121 shows the usage of the **vsrb** instruction. Each of the sixteen elements in the vectors, **vA**, and **vD**, is 8 bits long.



**Figure 6-121. vsrb—Shift Bits Right in Sixteen Integer Elements (8-Bit)**


# vsrh

# vsrh

Vector Shift Right Half Word

**vsrh**                                          **vD,vA,vB**

Form: VX



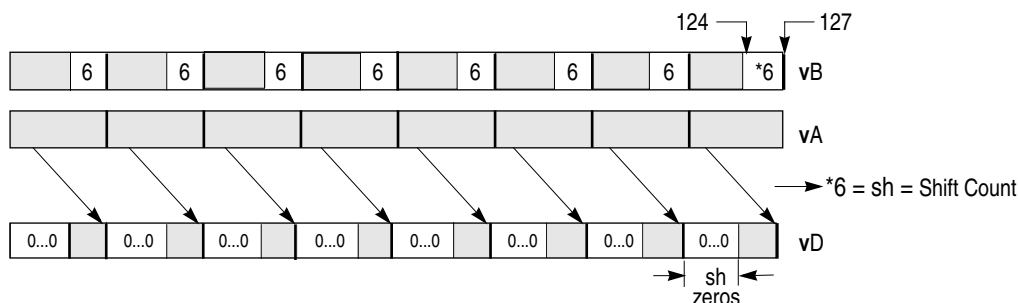
```
do i=0 to 127 by 16
    sh ← (vB)i+12:i+15
    vDi:i+15 ← (vA)i:i+15 >>ui sh
end
```

Each element is a half word. Each element in **vA** is shifted right by the number of bits specified in the low-order 4 bits of the corresponding element in **vB**. Bits shifted out of bit 15 of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-122 shows the usage of the **vsrh** instruction. Each of the eight elements in the vectors, **vA**, and **vD**, is 16 bits long.



**Figure 6-122. vsrh—Shift Bits Right for Eight Integer Elements (16-Bit)**

# vsro


Vector Shift Right Octet

# VSRO

vsro

vD,vA,vB

Form: VX



```
shb ← (vB)121:124
vD ← (vA) >>ui (shb || 0b000)
```

The contents of vA are shifted right by the number of bytes specified in vB[121–124]. Bytes shifted out of vA are lost. Zeros are supplied to the vacated bytes on the left. The result is placed into vD.

Other registers altered:

- None




Figure 6-123. vsro—Vector Shift Right Octet

# Vsrw


Vector Shift Right Word

# Vsrw

**vsrw**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
    sh ← (vB)i+(27):i+31
    vDi:i+31 ← (vA)i:i+31 >>ui sh
end
```

Each element is a word. Each element in **vA** is shifted right by the number of bits specified in the low-order 5 bits of the corresponding element in **vB**. Bits shifted out of bit 31 of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Figure 6-124 shows the usage of the **vsrw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-124. vsrw—Shift Bits Right in Four Integer Elements (32-Bit)**



# vsubcuw


Vector Subtract Carryout Unsigned Word

# vsubcuw

**vsubcuw**

**vD,vA,vB**

Form: VX




```
do i=0 to 127 by 32
    aop0:32 ← ZeroExtend((vA)i:i+31, 33)
    bop0:32 ← ZeroExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int -bop0:32 +int 1
    vDi:i+31 ← ZeroExtend(temp0, 32)
end
```

Each unsigned-integer word element in **vB** is subtracted from the corresponding unsigned-integer word element in **vA**. The complement of the borrow out of bit 0 of the 32-bit difference is zero-extended to 32 bits and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Figure 6-125 shows the usage of the **vsubcuw** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-125. vsubcuw—Subtract Carryout of Four Unsigned Integer Elements (32-Bit)**

# vsubfp


Vector Subtract Floating Point

# vsubfp

**vsubfp**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← RndToNearFP32((vA)i:i+31 -fp (vB)i:i+31)
end
```


Each single-precision floating-point word element in **vB** is subtracted from the corresponding single-precision floating-point word element in **vA**. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

If VSCR[NJ] = 1, every denormalized operand element is truncated to a 0 of the same sign before the operation is carried out, and each denormalized result element truncates to a 0 of the same sign.

Other registers altered:

- None

Figure 6-126 shows the usage of the **vsubfp** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-126. vsubfp—Subtract Four Floating Point Elements (32-Bit)**

# vsubsbs


Vector Subtract Signed Byte Saturate

# vsubsbs

**vsubsbs**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    aop0:8 ← SignExtend((vA)i:i+7, 9)
    bop0:8 ← SignExtend((vB)i:i+7, 9)
    temp0:8 ← aop0:8 +int -bop0:8 +int 1
    vDi:i+7 ← SItoSIsat(temp0:8, 8)
end
```

Each element is a byte. Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.


If the intermediate result is greater than  $(2^7-1)$  it saturates to  $(2^7-1)$  and if it is less than  $-2^7$  it saturates to  $-2^7$ , where 8 is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-127 shows the usage of the **vsubsbs** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-127. vsubsbs—Subtract Sixteen Signed Integer Elements (8-Bit)**

# vsubshs


Vector Subtract Signed Half Word Saturate

# vsubshs

**vsubshs**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
    aop0:16 ← SignExtend((vA)i:i+15,17)
    bop0:16 ← SignExtend((vB)i:i+15,17)
    temp0:16 ← aop0:16 +int -bop0:16 +int 1
    vDi:i+15 ← SItoSIsat(temp0:16,16)
end
```

Each element is a half word. Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.


If the intermediate result is greater than  $(2^{15}-1)$  it saturates to  $(2^{15}-1)$  and if it is less than  $-2^{15}$  it saturates to  $-2^{15}$ , where 16 is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-128 shows the usage of the **vsubshs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-128. vsubshs—Subtract Eight Signed Integer Elements (16-Bit)**

# vsubsws


Vector Subtract Signed Word Saturate

# vsubsws

**vsubsws**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    aop0:32 ← SignExtend((vA)i:i+31, 33)
    bop0:32 ← SignExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int -bop0:32 +int 1
    vDi:i+31 ← SItoSIsat(temp0:32, 32)
end
```

Each element is a word. Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.


If the intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ , where 32 is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-129 shows the usage of the **vsubsws** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-129. vsubsws—Subtract Four Signed Integer Elements (32-Bit)**

# vsububm


Vector Subtract Unsigned Byte Modulo

# vsububm

**vsububm**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    vDi:i+7 ← (vA)i:i+7 +int -(vB)i:i+7
end
```

Each element of **vsububm** is a byte.


Each integer element in **vB** is subtracted from the corresponding integer element in **vA**. The integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Note the **vsububm** instruction can be used for unsigned or signed integers.

Figure 6-130 shows the usage of the **vsububm** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-130. vsububm—Subtract Sixteen Integer Elements (8-Bit)**


# vsububs

# vsububs

Vector Subtract Unsigned Byte Saturate

**vsububs**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 8
    aop0:8 ← ZeroExtend((vA)i:i+7, 9)
    bop0:8 ← ZeroExtend((vB)i:i+7, 9)
    temp0:8 ← aop0:8 +int -bop0:8 +int 1
    vDi:i+7 ← SItOUIsat(temp0:8, 8)
end
```


Each element is a byte. Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where 8 is the length of the element. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-131 shows the usage of the **vsububs** instruction. Each of the sixteen elements in the vectors, **vA**, **vB**, and **vD**, is 8 bits long.



**Figure 6-131. vsububs—Subtract Sixteen Unsigned Integer Elements (8-Bit)**


# vsubuhm

# vsubuhm

Vector Subtract Signed Half Word Modulo

**vsubuhm**                      **vD,vA,vB**

Form: VX



```
do i=0 to 127 by 16
    vDi:i+15 ← (vA)i:i+15 +int -(vB)i:i+15
end
```


Each element is a half word. Each integer element in **vB** is subtracted from the corresponding integer element in **vA**. The integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Note the **vsubuhm** instruction can be used for unsigned or signed integers.

Figure 6-132 shows the usage of the **vsubuhm** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-132. vsubuhm—Subtract Eight Integer Elements (16-Bit)**



# vsubuhs


Vector Subtract Signed Half Word Saturate

# vsubuhs

**vsubuhs**

**vD,vA,vB**

Form: VX



```

do i=0 to 127 by 16
    aop0:16 ← ZeroExtend((vA)i:i+15,17)
    bop0:16 ← ZeroExtend((vB)i:i+n:1,17)
    temp0:16 ← aop0:n +int -bop0:16 +int 1
    vDi:i+15 ← SItouIsat(temp0:16,16)
end
    
```


Each element is a half word. Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where 16 is the length of the element. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-133 shows the usage of the **vsubuhs** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-133. vsubuhs—Subtract Eight Signed Integer Elements (16-Bit)**

# vsubuwm


# vsubuwm

Vector Subtract Unsigned Word Modulo

**vsubuwm**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    vDi:i+31 ← (vA)i:i+31 +int -(vB)i:i+31
end
```

Each element of **vsubuwm** is a word.


Each integer element in **vB** is subtracted from the corresponding integer element in **vA**. The integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Note the **vsubuwm** instruction can be used for unsigned or signed integers.

Figure 6-134 shows the usage of the **vsubuwm** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-134. vsubuwm—Subtract Four Integer Elements (32-Bit)**

# vsubuws


Vector Subtract Unsigned Word Saturate

# vsubuws

**vsubuws**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    aop0:32 ← ZeroExtend((vA)i:i+31, 33)
    bop0:32 ← ZeroExtend((vB)i:i+31, 33)
    temp0:32 ← aop0:32 +int -bop0:32 +int 1
    vDi:i+31 ← SItouIsat(temp0:32, 32)
end
```


Each element is a word. Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where 32 is the length of the element. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT

Figure 6-135 shows the usage of the **vsubuws** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-135. vsubuws—Subtract Four Signed Integer Elements (32-Bit)**

# vsumsws


Vector Sum Across Signed Word Saturate

# vsumsws

**vsumsws**

**vD,vA,vB**

Form: VX



```


temp0:34 ← SignExtend((vB)96:127, 35)
do i=0 to 127 by 32
    temp0:34 ← temp0:34 +int SignExtend((vA)i:i+31, 35)
    vD ← 960 || SItoSIsat(temp0:34, 32)
end
  
```

The signed-integer sum of the four signed-integer word elements in **vA** is added to the signed-integer word element in bits of **vB**[96-127]. If the intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . The signed-integer result is placed into bits **vD**[96-127]. Bits **vD**[0-95] are cleared.

Other registers altered:

- SAT

Figure 6-136 shows the usage of the **vsumsws** instruction. Each of the four elements in the vectors, **vA**, **vB**, and **vD**, is 32 bits long.



**Figure 6-136. vsumsws—Sum Four Signed Integer Elements (32-Bit)**

# vsum2sws


# vsum2sws

Vector Sum Across Partial (1/2) Signed Word Saturate

vsum2sws

vD,vA,vB

Form: VX



```

do i=0 to 127 by 64
    temp0:33 ← SignExtend((vB)i+32:i+63,34)
    do j=0 to 63 by 32
        temp0:33 ← temp0:33 +int SignExtend((vA)i+j:i+j+31,34)
    end
    vDi:i+63 ← 320 || SItoSIsat(temp0:33,32)
end
    
```

The signed-integer sum of the first two signed-integer word elements in register vA is added to the signed-integer word element in vB[32–63]. If the intermediate result is greater than (2<sup>31</sup>-1) it saturates to (2<sup>31</sup>-1) and if it is less than -2<sup>31</sup> it saturates to -2<sup>31</sup>. The signed-integer result is placed into vD[32–63]. The signed-integer sum of the last two signed-integer word elements in register vA is added to the signed-integer word element in vB[96-127]. If the intermediate result is greater than (2<sup>31</sup>-1) it saturates to (2<sup>31</sup>-1) and if it is less than -2<sup>31</sup> it saturates to -2<sup>31</sup>. The signed-integer result is placed into vD[96–127]. The register vD[0–31,64–95] are cleared to 0.

Other registers altered:

- SAT

Figure 6-137 shows the usage of the vsum2sws instruction. Each of the four elements in the vectors, vA, vB, and vD, is 32 bits long.




Figure 6-137. vsum2sws—Two Sums in the Four Signed Integer Elements (32-Bit)

# vsum4sbs


# vsum4sbs

Vector Sum Across Partial (1/4) Signed Byte Saturate

**vsum4sbs**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    temp0:32 ← SignExtend((vB)i:i+31,33)
    do j=0 to 31 by 8
        temp0:32 ← temp0:32 +int SignExtend((vA)i+j:i+j+7,33)
    end
    vDi:i+31 ← SItoSIsat(temp0:32,32)
end
```


For each word element in **vB** the following operations are performed in the order shown.

- The signed-integer sum of the four signed-integer byte elements contained in the corresponding word element of register **vA** is added to the signed-integer word element in register **vB**.
- If the intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT

Figure 6-138 shows the usage of the **vsum4sbs** instruction. Each of the sixteen elements in the vector **vA**, is 8 bits long. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-138. vsum4sbs—Four Sums in the Integer Elements (32-Bit)**

# vsum4shs


# vsum4shs

Vector Sum Across Partial (1/4) Signed Half Word Saturate

vsum4shs

vD,vA,vB

Form: VX



```

do i=0 to 127 by 32
    temp0:32 ← SignExtend((vB)i:i+31,33)
    do j=0 to 31 by 16
        temp0:32 ← temp0:32 +int SignExtend((vA)i+j:i+j+15,33)
    end
    vDi:i+31 ← SItoSIsat(temp0:32,32)
end
    
```

For each word element in register vB the following operations are performed, in the order shown.

- The signed-integer sum of the two signed-integer halfword elements contained in the corresponding word element of register vA is added to the signed-integer word element in vB.
- If the intermediate result is greater than  $(2^{31}-1)$  it saturates to  $(2^{31}-1)$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of vD.

Other registers altered:

- SAT

Figure 6-139 shows the usage of the vsum4shs instruction. Each of the eight elements in the vector vA, is 16 bits long. Each of the four elements in the vectors vB and vD is 32 bits long.




Figure 6-139. vsum4shs—Four Sums in the Integer Elements (32-Bit)

# vsum4ubs


# vsum4ubs

Vector Sum Across Partial (1/4) Unsigned Byte Saturate

**vsum4ubs**

**vD,vA,vB**

Form: VX



```
do i=0 to 127 by 32
    temp0:32 ← ZeroExtend((vB)i:i+31, 33)
    do j=0 to 31 by 8
        temp0:32 ← temp0:32 +int ZeroExtend((vA)i+j:i+j+7, 33)
    end
    vDi:i+31 ← UItoUISat(temp0:32, 32)
end
```


For each word element in **vB** the following operations are performed in the order shown.

- The unsigned-integer sum of the four unsigned-integer byte elements contained in the corresponding word element of register **vA** is added to the unsigned-integer word element in register **vB**.
- If the intermediate result is greater than  $(2^{32}-1)$  it saturates to  $(2^{32}-1)$ .
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT

Figure 6-140 shows the usage of the **vsum4ubs** instruction. Each of the four elements in the vector **vA**, is 8 bits long. Each of the four elements in the vectors **vB** and **vD** is 32 bits long.



**Figure 6-140. vsum4ubs—Four Sums in the Integer Elements (32-Bit)**



# vupkhp<sub>x</sub>


Vector Unpack High Pixel16

# vupkhp<sub>x</sub>

vupkhp<sub>x</sub>

vD,vB

Form: VX



do i=0 to 63 by 16

```

vDi*2:(i*2)+7 ← SignExtend((vB)i, 8)
vD(i*2)+8:(i*2)+15 ← ZeroExtend((vB)i+1:i+5, 8)
vD(i*2)+16:(i*2)+23 ← ZeroExtend((vB)i+6:i+10, 8)
vD(i*2)+24:(i*2)+31 ← ZeroExtend((vB)i+11:i+15, 8)
    
```

end

Each halfword element in the high-order half of register vB is unpacked to produce a 32-bit value as described below and placed, in the same order, into the four words of vD.

A halfword is unpacked to 32 bits by concatenating, in order, the results of the following operations.


- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1–5 of the halfword to 8 bits
- zero-extend bits 6–10 of the halfword to 8 bits
- zero-extend bits 11–15 of the halfword to 8 bits

Other registers altered:

- None

The source and target elements can be considered to be 16-bit and 32-bit "pixels" respectively, having the formats described in the programming note for the Vector Pack Pixel instruction.

Figure 6-141 shows the usage of the vupkhp<sub>x</sub> instruction. Each of the eight elements in the vectors, vB, is 16 bits long. Each of the four elements in the vectors, vD, is 32 bits long.



**Figure 6-141. vupkhp<sub>x</sub>—Unpack High-Order Elements (16 bit) to Elements (32-Bit)**

# vupkhsb


Vector Unpack High Signed Byte

# vupkhsb

**vupkhsb**

**vD,vB**

Form: VX




```
do i=0 to 63 by 8
    vDi*2:(i*2)+15 ← SignExtend((vB)i:i+7,16)
end
```

Each signed integer byte element in the high-order half of register **vB** is sign-extended to produce a 16-bit signed integer and placed, in the same order, into the eight halfwords of register **vD**.

Other registers altered:

- None

Figure 6-142 shows the usage of the **vupkhsb** instruction. Each of the sixteen elements in the vectors, **vB**, is 8 bits long. Each of the eight elements in the vectors, **vD**, is 16 bits long.



**Figure 6-142. vupkhsb—Unpack High-Order Signed Integer Elements (8-Bit) to Signed Integer Elements (16-Bit)**

# vupkhsh


Vector Unpack High Signed Half Word

# vupkhsh

**vupkhsh**

**vD,vB**

Form: VX




```
do i=0 to 63 by 16
    vDi*2:(i*2)+31 ← SignExtend((vB)i:i+15,32)
end
```

Each signed integer halfword element in the high-order half of register **vB** is sign-extended to produce a 32-bit signed integer and placed, in the same order, into the four words of register **vD**.

Other registers altered:

- None

Figure 6-143 shows the usage of the **vupkhsh** instruction. Each of the eight elements in the vectors **vB** and **vD** is 16 bits long.



**Figure 6-143. vupkhsh—Unpack Signed Integer Elements (16-Bit) to Signed Integer Elements (32-Bit)**

# vupklpx


Vector Unpack Low Pixel16

# vupklpx

**vupklpx**

**vD,vB**

Form: VX



```
do i=0 to 63 by 16
    vDi*2:(i*2)+7 ← SignExtend((vB)i+64, 8)
    vD(i*2)+8:(i*2)+15 ← ZeroExtend((vB)i+65:i+69, 8)
    vD(i*2)+16:(i*2)+23 ← ZeroExtend((vB)i+70:i+74, 8)
    vD(i*2)+24:(i*2)+31 ← ZeroExtend((vB)i+75:i+79, 8)
end
```

Each halfword element in the low-order half of register **vB** is unpacked to produce a 32-bit value as described below and placed, in the same order, into the four words of register **vD**.

A halfword is unpacked to 32 bits by concatenating, in order, the results of the following operations.


- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1–5 of the halfword to 8 bits
- zero-extend bits 6–10 of the halfword to 8 bits
- zero-extend bits 11–15 of the halfword to 8 bits

Other registers altered:

- None

Programming note: Notice that the unpacking done by the Vector Unpack Pixel instructions does not reverse the packing done by the Vector Pack Pixel instruction. Specifically, if a 16-bit pixel is unpacked to a 32-bit pixel which is then packed to a 16-bit pixel, the resulting 16-bit pixel will not, in general, be equal to the original 16-bit pixel (because, for each channel except the first, Vector Unpack Pixel inserts high-order bits while Vector Pack Pixel discards low-order bits).

Figure 6-144 shows the usage of the **vupklpx** instruction. Each of the eight elements in the vectors, **vB**, is 16 bits long. Each of the four elements in the vectors, **vD**, is 32 bits long.



**Figure 6-144. vupklpx—Unpack Low-order Elements (16-Bit) to Elements (32-Bit)**

# vupklsb


Vector Unpack Low Signed Byte

# vupklsb

**vupklsb**

**vD,vB**

Form: VX




```
do i=0 to 63 by 8
    vDi*2:(i*2)+15 ← SignExtend((vB)i+64:i+71,16)
end
```

Each signed integer byte element in the low-order half of register **vB** is sign-extended to produce a 16-bit signed integer and placed, in the same order, into the eight halfwords of register **vD**.

Other registers altered:

- None

Figure 6-145 shows the usage of the **vaddubs** instruction. Each of the sixteen elements in the vectors **vB** and **vD** is 8 bits long.



**Figure 6-145. vupklsb—Unpack Low-Order Elements (8-Bit) to Elements (16-Bit)**

# vupklsh


Vector Unpack Low Signed Half Word

# vupklsh

**vupklsh**

**vD,vB**

Form: VX




```
do i=0 to 63 by 16
    vDi*2:(i*2)+31 ← SignExtend((vB)i+64:i+79, 32)
end
```

Each signed integer half word element in the low-order half of register **vB** is sign-extended to produce a 32-bit signed integer and placed, in the same order, into the four words of register **vD**.

Other registers altered:

- None

Figure 6-146 shows the usage of the **vupklpx** instruction. Each of the eight elements in the vectors, **vA**, **vB**, and **vD**, is 16 bits long.



**Figure 6-146. vupklsh—Unpack Low-Order Signed Integer Elements (16-Bit) to Signed Integer Elements (32-Bit)**

# VXor


Vector Logical XOR

# VXOR

**vxor**

**vD,vA,vB**

Form: VX




$$vD \leftarrow (vA) \oplus (vB)$$

The contents of vA are XORed with the contents of register vB and the result is placed into register vD.

Other registers altered:

- None

Figure 6-147 shows the usage of the **vxor** instruction.



**Figure 6-147. vxor—Bitwise XOR (128-Bit)**





# Appendix A

## AltiVec Instruction Set Listings

This appendix lists the instruction set for AltiVec™ technology. Instructions are sorted by mnemonic, opcode, and form. Also included in this appendix is a quick reference table that contains general information, such as the architecture level, privilege level, and form, and indicates if the instruction is optional.

Note that split fields, which represent the concatenation of sequences from left to right, are shown in lower case.

### A.1 Instructions Sorted by Mnemonic in Decimal Format

Table A-1 lists the instructions implemented in the AltiVec architecture in alphabetical order by mnemonic. The primary and extended opcodes are decimal numbers.

Key:

 Reserved bits

**Table A-1. Instruction Sorted by Mnemonic in Decimal Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>dss</b>	31	0	0_0	STRM	0_0000	0000_0															822							0
<b>dssall</b>	31	1	0_0	STRM	0_0000	0000_0															822							0
<b>dst</b>	31	0	0_0	STRM	A	B															342							0
<b>dstst</b>	31	0	0_0	STRM	A	B															374							0
<b>dststt</b>	31	1	0_0	STRM	A	B															374							0
<b>dstt</b>	31	1	0_0	STRM	A	B															342							0
<b>lvebx</b>	31		vD		A	B															7							0
<b>lvehx</b>	31		vD		A	B															39							0
<b>lvewx</b>	31		vD		A	B															71							0
<b>lvsl</b>	31		vD		A	B															6							0
<b>lvslr</b>	31		vD		A	B															38							0
<b>lvx</b>	31		vD		A	B															103							0

**Table A-1. Instruction Sorted by Mnemonic in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
lvxl	31		vD				A				B				359						0							
mfvscr	04		vD				0_0000				0000_0				1540													
mtvscr	04		00_000				0_0000				vB				1604													
stvebx	31		vS				A				B				135						0							
stvehx	31		vS				A				B				167						0							
stviewx	31		vS				A				B				199						0							
stvx	31		vS				A				B				231						0							
stvxl	31		vS				A				B				487						0							
vaddcuw	04		vD				vA				vB				384													
vaddfp	04		vD				vA				vB				10													
vaddsbs	04		vD				vA				vB				768													
vaddshs	04		vD				vA				vB				832													
vaddsws	04		vD				vA				vB				896													
vaddubm	04		vD				vA				vB				0													
vaddubs	04		vD				vA				vB				512													
vadduhm	04		vD				vA				vB				64													
vadduhs	04		vD				vA				vB				576													
vadduwm	04		vD				vA				vB				128													
vadduws	04		vD				vA				vB				640													
vand	04		vD				vA				vB				1028													
vandc	04		vD				vA				vB				1092													
vavgsb	04		vD				vA				vB				1282													
vavgsh	04		vD				vA				vB				1346													
vavgsw	04		vD				vA				vB				1410													
vavgub	04		vD				vA				vB				1026													
vavguh	04		vD				vA				vB				1090													
vavguw	04		vD				vA				vB				1154													
vcfsx	04		vD				UIMM				vB				842													
vcfux	04		vD				UIMM				vB				778													
vcmpbfp	04		vD				vA				vB			Rc	966													
vcmpeqfp	04		vD				vA				vB			Rc	198													
vcmpequb	04		vD				vA				vB			Rc	6													
vcmpequh	04		vD				vA				vB			Rc	70													

**Table A-1. Instruction Sorted by Mnemonic in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vcmpequwx	04				vD					vA					vB			Rc										134
vcmpgefp	04				vD					vA					vB			Rc										454
vcmpgtfp	04				vD					vA					vB			Rc										710
vcmpgtsb	04				vD					vA					vB			Rc										774
vcmpgtsh	04				vD					vA					vB			Rc										838
vcmpgtsw	04				vD					vA					vB			Rc										902
vcmpgtub	04				vD					vA					vB			Rc										518
vcmpgtuh	04				vD					vA					vB			Rc										582
vcmpgtuw	04				vD					vA					vB			Rc										646
vctsxs	04				vD					UIMM					vB													970
vctuxs	04				vD					UIMM					vB													906
vexptefp	04				vD					0_0000					vB													394
vlogefp	04				vD					0_0000					vB													458
vmaddfp	04				vD					vA					vB				vC									46
vmaxfp	04				vD					vA					vB													1034
vmaxsb	04				vD					vA					vB													258
vmaxsh	04				vD					vA					vB													322
vmaxsw	04				vD					vA					vB													386
vmaxub	04				vD					vA					vB													2
vmaxuh	04				vD					vA					vB													66
vmaxuw	04				vD					vA					vB													130
vmhaddshs	04				vD					vA					vB				vC									32
vmhraddshs	04				vD					vA					vB				vC									33
vminfp	04				vD					vA					vB													1098
vminsb	04				vD					vA					vB													770
vminsh	04				vD					vA					vB													834
vminsw	04				vD					vA					vB													898
vminub	04				vD					vA					vB													514
vminuh	04				vD					vA					vB													578
vminuw	04				vD					vA					vB													642
vmladduhm	04				vD					vA					vB				vC									34
vmrghb	04				vD					vA					vB													12
vmrghh	04				vD					vA					vB													76

**Table A-1. Instruction Sorted by Mnemonic in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vmrghw	04			vD					vA					vB															140
vmrglb	04			vD					vA					vB															268
vmrglh	04			vD					vA					vB															332
vmrglw	04			vD					vA					vB															396
vmsummbm	04			vD					vA					vB						vC									37
vmsumshm	04			vD					vA					vB						vC									40
vmsumshs	04			vD					vA					vB						vC									41
vmsumubm	04			vD					vA					vB						vC									36
vmsumuhm	04			vD					vA					vB						vC									38
vmsumuhs	04			vD					vA					vB						vC									39
vmulesb	04			vD					vA					vB															776
vmulesh	04			vD					vA					vB															840
vmuleub	04			vD					vA					vB															520
vmuleuh	04			vD					vA					vB															584
vmulosb	04			vD					vA					vB															264
vmulosh	04			vD					vA					vB															328
vmuloub	04			vD					vA					vB															8
vmulouh	04			vD					vA					vB															72
vnmsubfp	04			vD					vA					vB						vC									47
vnor	04			vD					vA					vB															1284
vor	04			vD					vA					vB															1156
vperm	04			vD					vA					vB						vC									43
vpkpx	04			vD					vA					vB															782
vpkshss	04			vD					vA					vB															398
vpkshus	04			vD					vA					vB															270
vpkswss	04			vD					vA					vB															462
vpkswus	04			vD					vA					vB															334
vpkuhum	04			vD					vA					vB															14
vpkuhus	04			vD					vA					vB															142
vpkuwum	04			vD					vA					vB															78
vpkuwus	04			vD					vA					vB															206
vrefp	04			vD					0_0000					vB															266
vrfim	04			vD					0_0000					vB															714

**Table A-1. Instruction Sorted by Mnemonic in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vrfn</b>	04				vD				0_0000					vB									522					
<b>vrfip</b>	04				vD				0_0000					vB									650					
<b>vrfiz</b>	04				vD				0_0000					vB									586					
<b>vrlb</b>	04				vD				vA					vB									4					
<b>vrlh</b>	04				vD				vA					vB									68					
<b>vrlw</b>	04				vD				vA					vB									132					
<b>vrsqrtefp</b>	04				vD				0_0000					vB									330					
<b>vsel</b>	04				vD				vA					vB						vC					42			
<b>vsl</b>	04				vD				vA					vB									452					
<b>vslb</b>	04				vD				vA					vB									260					
<b>vsldoi</b>	04				vD				vA					vB				0		SH					44			
<b>vslh</b>	04				vD				vA					vB									324					
<b>vslo</b>	04				vD				vA					vB									1036					
<b>vslw</b>	04				vD				vA					vB									388					
<b>vspltb</b>	04				vD				UIMM					vB									524					
<b>vsplth</b>	04				vD				UIMM					vB									588					
<b>vspltisb</b>	04				vD				SIMM					0000_0									780					
<b>vspltish</b>	04				vD				SIMM					0000_0									844					
<b>vspltisw</b>	04				vD				SIMM					0000_0									908					
<b>vspltw</b>	04				vD				UIMM					vB									652					
<b>vsr</b>	04				vD				vA					vB									708					
<b>vsrab</b>	04				vD				vA					vB									772					
<b>vsrah</b>	04				vD				vA					vB									836					
<b>vsraw</b>	04				vD				vA					vB									900					
<b>vsrb</b>	04				vD				vA					vB									516					
<b>vsrh</b>	04				vD				vA					vB									580					
<b>vsro</b>	04				vD				vA					vB									1100					
<b>vsrw</b>	04				vD				vA					vB									644					
<b>vsubcuw</b>	04				vD				vA					vB									1408					
<b>vsubfp</b>	04				vD				vA					vB									74					
<b>vsubsubs</b>	04				vD				vA					vB									1792					
<b>vsubshs</b>	04				vD				vA					vB									1856					
<b>vsubsws</b>	04				vD				vA					vB									1920					



# Appendix B

## Instructions Sorted by Mnemonic in Binary Format

### B.1 Instructions Sorted by Mnemonic in Binary Format

Table B-1 lists the instructions implemented in the AltiVec architecture in alphabetical order by mnemonic. The primary and extended opcodes are decimal numbers.

Key:

 Reserved bits

**Table B-1. Instructions Sorted by Mnemonic in Binary Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>dss</b>	0111_11	0	0_0	STRM	0_0000		0000_0		110_0110_110								0											
<b>dssall</b>	0111_11	1	0_0	STRM	0_0000		0000_0		110_0110_110								0											
<b>dst</b>	0111_11	0	0_0	STRM	A		B		010_1010_110								0											
<b>dstst</b>	0111_11	0	0_0	STRM	A		B		010_1110_110								0											
<b>dststt</b>	0111_11	1	0_0	STRM	A		B		001_1110_110								0											
<b>dstt</b>	0111_11	1	0_0	STRM	A		B		010_1010_110								0											
<b>lvebx</b>	0111_11	vD		A		B		000_0000_111								0												
<b>lvehx</b>	0111_11	vD		A		B		000_0100_111								0												
<b>lviewx</b>	0111_11	vD		A		B		000_1000_111								0												
<b>lvsl</b>	0111_11	vD		A		B		000_0000_110								0												
<b>lvslr</b>	0111_11	vD		A		B		000_0100_110								0												
<b>lvx</b>	0111_11	vD		A		B		000_1100_111								0												
<b>lvxl</b>	0111_11	vD		A		B		010_1100_111								0												
<b>mfvscr</b>	0001_00	vD		0_0000		0000_0		110_0000_0100								0												
<b>mtvscr</b>	0001_00	00_000		0_0000		vB		110_0100_0100								0												
<b>stvebx</b>	0111_11	vS		A		B		001_0000_111								0												
<b>stvehx</b>	0111_11	vS		A		B		001_0100_111								0												

**Table B-1. Instructions Sorted by Mnemonic in Binary Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
<b>stviewx</b>	0111_11				vS					A					B						001_1000_111								0
<b>stvx</b>	0111_11				vS					A					B						001_1100_111								0
<b>stvxl</b>	0111_11				vS					A					B						011_1100_111								0
<b>vaddcuw</b>	0001_00				vD					vA					vB						001_1000_0000								
<b>vaddfp</b>	0001_00				vD					vA					vB						000_0000_1010								
<b>vaddsbs</b>	0001_00				vD					vA					vB						011_0000_0000								
<b>vaddshs</b>	0001_00				vD					vA					vB						011_0100_0000								
<b>vaddsws</b>	0001_00				vD					vA					vB						011_1000_0000								
<b>vaddubm</b>	0001_00				vD					vA					vB						000_0000_0000								
<b>vaddubs</b>	0001_00				vD					vA					vB						010_0000_0000								
<b>vadduhm</b>	0001_00				vD					vA					vB						000_0100_0000								
<b>vadduhs</b>	0001_00				vD					vA					vB						010_0100_0000								
<b>vadduwm</b>	0001_00				vD					vA					vB						000_1000_0000								
<b>vadduws</b>	0001_00				vD					vA					vB						010_1000_0000								
<b>vand</b>	0001_00				vD					vA					vB						100_0000_0100								
<b>vandc</b>	0001_00				vD					vA					vB						100_0100_0100								
<b>vavgsb</b>	0001_00				vD					vA					vB						101_0000_0010								
<b>vavgsh</b>	0001_00				vD					vA					vB						101_0100_0010								
<b>vavgsw</b>	0001_00				vD					vA					vB						101_1000_0010								
<b>vavgub</b>	0001_00				vD					vA					vB						100_0000_0010								
<b>vavguh</b>	0001_00				vD					vA					vB						100_0100_0010								
<b>vavguw</b>	0001_00				vD					vA					vB						100_1000_0010								
<b>vcfsx</b>	0001_00				vD					UIMM					vB						011_0100_1010								
<b>vcfux</b>	0001_00				vD					UIMM					vB						011_0000_1010								
<b>vcmpbfp</b>	0001_00				vD					vA					vB		Rc				11_1100_0110								
<b>vcmpeqfp</b>	0001_00				vD					vA					vB		Rc				00_1100_0110								
<b>vcmpequb</b>	0001_00				vD					vA					vB		Rc				00_0000_0110								
<b>vcmpequh</b>	0001_00				vD					vA					vB		Rc				00_0100_0110								
<b>vcmpequw</b>	0001_00				vD					vA					vB		Rc				00_1000_0110								
<b>vcmpgef</b>	0001_00				vD					vA					vB		Rc				01_1100_0110								
<b>vcmpgtf</b>	0001_00				vD					vA					vB		Rc				10_1100_0110								
<b>vcmpgtsb</b>	0001_00				vD					vA					vB		Rc				11_0000_0110								
<b>vcmpgtsh</b>	0001_00				vD					vA					vB		Rc				11_0100_0110								



### Table B-1. Instructions Sorted by Mnemonic in Binary Format

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vcmpgtswx</b>	0001_00		vD		vA		vB		Rc	11_1000_0110																		
<b>vcmpgtubx</b>	0001_00		vD		vA		vB		Rc	10_0000_0110																		
<b>vcmpgtuhx</b>	0001_00		vD		vA		vB		Rc	10_0100_0110																		
<b>vcmpgtuwx</b>	0001_00		vD		vA		vB		Rc	10_1000_0110																		
<b>vctxsx</b>	0001_00		vD		UIMM		vB		011_1100_1010																			
<b>vctuxs</b>	0001_00		vD		UIMM		vB		011_1000_1010																			
<b>vexptefp</b>	0001_00		vD		0_0000		vB		001_1000_1010																			
<b>vlogefp</b>	0001_00		vD		0_0000		vB		001_1100_1010																			
<b>vmaddfp</b>	0001_00		vD		vA		vB		vC	10_1110																		
<b>vmaxfp</b>	0001_00		vD		vA		vB		100_0000_1010																			
<b>vmaxsb</b>	0001_00		vD		vA		vB		001_0000_0010																			
<b>vmaxsh</b>	0001_00		vD		vA		vB		001_0100_0010																			
<b>vmaxsw</b>	0001_00		vD		vA		vB		001_1000_0010																			
<b>vmaxub</b>	0001_00		vD		vA		vB		0000_0000_0010																			
<b>vmaxuh</b>	0001_00		vD		vA		vB		0100_0010																			
<b>vmaxuw</b>	0001_00		vD		vA		vB		1000_0010																			
<b>vmhaddshs</b>	0001_00		vD		vA		vB		vC	10_0000																		
<b>vmhraddshs</b>	0001_00		vD		vA		vB		vC	10_0001																		
<b>vminfp</b>	0001_00		vD		vA		vB		100_0100_1010																			
<b>vminsb</b>	0001_00		vD		vA		vB		011_0000_0010																			
<b>vminsh</b>	0001_00		vD		vA		vB		011_0100_0010																			
<b>vminsw</b>	0001_00		vD		vA		vB		011_1000_0010																			
<b>vminub</b>	0001_00		vD		vA		vB		010_0000_0010																			
<b>vminuh</b>	0001_00		vD		vA		vB		010_0100_0010																			
<b>vminuw</b>	0001_00		vD		vA		vB		010_1000_0010																			
<b>vmladduhm</b>	0001_00		vD		vA		vB		vC	10_0010																		
<b>vmrghb</b>	0001_00		vD		vA		vB		000_0000_1100																			
<b>vmrghh</b>	0001_00		vD		vA		vB		000_0100_1100																			
<b>vmrghw</b>	0001_00		vD		vA		vB		000_1000_1100																			
<b>vmrglb</b>	0001_00		vD		vA		vB		001_0000_1100																			
<b>vmrglh</b>	0001_00		vD		vA		vB		001_0100_1100																			
<b>vmrglw</b>	0001_00		vD		vA		vB		001_1000_1100																			
<b>vmsummbm</b>	0001_00		vD		vA		vB		vC	10_0101																		

**Table B-1. Instructions Sorted by Mnemonic in Binary Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vmsumshm	0001_00				vD					vA					vB					vC								10_1000
vmsumshs	0001_00				vD					vA					vB					vC								10_1001
vmsumubm	0001_00				vD					vA					vB					vC								10_0100
vmsumuhm	0001_00				vD					vA					vB					vC								10_0110
vmsumuhs	0001_00				vD					vA					vB					vC								10_0111
vmulesb	0001_00				vD					vA					vB													011_0000_1000
vmulesh	0001_00				vD					vA					vB													011_0100_1000
vmuleub	0001_00				vD					vA					vB													010_0000_1000
vmuleuh	0001_00				vD					vA					vB													010_0100_1000
vmulosb	0001_00				vD					vA					vB													001_0000_1000
vmulosh	0001_00				vD					vA					vB													001_0100_1000
vmuloub	0001_00				vD					vA					vB													000_0000_1000
vmulouh	0001_00				vD					vA					vB													000_0100_1000
vnmsubfp	0001_00				vD					vA					vB					vC								10_1111
vnor	0001_00				vD					vA					vB													101_0000_0100
vor	0001_00				vD					vA					vB													100_1000_0100
vperm	0001_00				vD					vA					vB					vC								10_1011
vpkpx	0001_00				vD					vA					vB													011_0000_1110
vpkshss	0001_00				vD					vA					vB													001_1000_1110
vpkshus	0001_00				vD					vA					vB													001_0000_1110
vpkswss	0001_00				vD					vA					vB													001_1100_1110
vpkswus	0001_00				vD					vA					vB													001_0100_1110
vpkuhum	0001_00				vD					vA					vB													000_0000_1110
vpkuhus	0001_00				vD					vA					vB													000_1000_1110
vpkuwum	0001_00				vD					vA					vB													000_100_1110
vpkuwus	0001_00				vD					vA					vB													000_1100_1110
vrefp	0001_00				vD					0_0000					vB													001_0000_1010
vrfim	0001_00				vD					0_0000					vB													010_1100_1010
vrfin	0001_00				vD					0_0000					vB													010_0000_1010
vrfip	0001_00				vD					0_0000					vB													010_1000_1010
vrfiz	0001_00				vD					0_0000					vB													010_0100_1010
vrlb	0001_00				vD					vA					vB													000_0000_0100
vrlh	0001_00				vD					vA					vB													000_0100_0100

### Table B-1. Instructions Sorted by Mnemonic in Binary Format

Name	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
<b>vrlw</b>	0001_00					vD					vA					vB					000_1000_0100																
<b>vsqrtefp</b>	0001_00					vD					0_0000					vB					001_0100_1010																
<b>vsel</b>	0001_00					vD					vA					vB					vC					10_1010											
<b>vsl</b>	0001_00					vD					vA					vB					1_1100_0100																
<b>vs1b</b>	0001_00					vD					vA					vB					1_0000_0100																
<b>vsldoi</b>	0001_00					vD					vA					vB					0	SH					10_1100										
<b>vslh</b>	0001_00					vD					vA					vB					01_0100_0100																
<b>vslo</b>	0001_00					vD					vA					vB					100_0000_1100																
<b>vslw</b>	0001_00					vD					vA					vB					001_1000_0100																
<b>vspltb</b>	0001_00					vD					UIMM					vB					010_0000_1100																
<b>vsplth</b>	0001_00					vD					UIMM					vB					010_0100_1100																
<b>vspltisb</b>	0001_00					vD					SIMM					0000_0					011_0000_1100																
<b>vspltish</b>	0001_00					vD					SIMM					0000_0					011_0100_1100																
<b>vspltisw</b>	0001_00					vD					SIMM					0000_0					011_1000_1100																
<b>vspltw</b>	0001_00					vD					UIMM					vB					010_1000_1100																
<b>vsr</b>	0001_00					vD					vA					vB					010_1100_0100																
<b>vsrab</b>	0001_00					vD					vA					vB					011_0000_0100																
<b>vsrah</b>	0001_00					vD					vA					vB					011_0100_0100																
<b>vsraw</b>	0001_00					vD					vA					vB					011_1000_0100																
<b>vsrb</b>	0001_00					vD					vA					vB					010_0000_0100																
<b>vsrh</b>	0001_00					vD					vA					vB					010_0100_0100																
<b>vsro</b>	0001_00					vD					vA					vB					100_0100_1100																
<b>vsrw</b>	0001_00					vD					vA					vB					010_1000_0100																
<b>vsubcuw</b>	0001_00					vD					vA					vB					101_1000_0000																
<b>vsubfp</b>	0001_00					vD					vA					vB					000_0100_1010																
<b>vsubsubs</b>	0001_00					vD					vA					vB					111_0000_0000																
<b>vsubshs</b>	0001_00					vD					vA					vB					111_0100_0000																
<b>vsubsws</b>	0001_00					vD					vA					vB					111_1000_0000																
<b>vsububm</b>	0001_00					vD					vA					vB					100_0000_0000																
<b>vsububs</b>	0001_00					vD					vA					vB					110_0000_0000																
<b>vsubuhm</b>	0001_00					vD					vA					vB					100_0100_0000																
<b>vsubuhs</b>	0001_00					vD					vA					vB					110_0100_0000																
<b>vsubuwm</b>	0001_00					vD					vA					vB					100_1000_0000																

**Table B-1. Instructions Sorted by Mnemonic in Binary Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
<b>vsubuws</b>	0001_00			vD					vA					vB								110_1000_0000							
<b>vsumsws</b>	0001_00			vD					vA					vB								111_1000_1000							
<b>vsum2sws</b>	0001_00			vD					vA					vB								110_1000_1000							
<b>vsum4sbs</b>	0001_00			vD					vA					vB								111_0000_1000							
<b>vsum4shs</b>	0001_00			vD					vA					vB								110_0100_1000							
<b>vsum4ubs</b>	0001_00			vD					vA					vB								110_0000_1000							
<b>vupkhpX</b>	0001_00			vD					0_0000					vB								011_0100_1110							
<b>vupkhsb</b>	0001_00			vD					0_0000					vB								010_0000_1110							
<b>vupkhsh</b>	0001_00			vD					0_0000					vB								010_0100_1110							
<b>vupklpx</b>	0001_00			vD					0_0000					vB								011_1100_1110							
<b>vupklsb</b>	0001_00			vD					0_0000					vB								010_1000_1110							
<b>vupklsh</b>	0001_00			vD					0_0000					vB								010_1100_1110							
<b>vxor</b>	0001_00			vD					vA					vB								100_1100_0100							

# Appendix C

## Instructions Sorted by Opcode

### C.1 Instructions Sorted by Opcode in Decimal Format

Table C-1 lists AltiVec instructions grouped by opcode in decimal format.

Key:

 Reserved bits

Table C-1. Instructions Sorted by Opcode in Decimal Format

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vmhaddshs	04				vD					vA					vB					vC								32
vmhraddshs	04				vD					vA					vB					vC								33
vmladduhm	04				vD					vA					vB					vC								34
vmsumubm	04				vD					vA					vB					vC								36
vmsummbm	04				vD					vA					vB					vC								37
vmsumuhm	04				vD					vA					vB					vC								38
vmsumuhs	04				vD					vA					vB					vC								39
vmsumshm	04				vD					vA					vB					vC								40
vmsumshs	04				vD					vA					vB					vC								41
vsel	04				vD					vA					vB					vC								42
vperm	04				vD					vA					vB					vC								43
vsldoi	04				vD					vA					vB			0		SH								44
vmaddfp	04				vD					vA					vB													46
vnmsubfp	04				vD					vA					vB					vC								47
vaddubm	04				vD					vA					vB													0
vadduhm	04				vD					vA					vB													64
vadduwm	04				vD					vA					vB													128
vaddcuw	04				vD					vA					vB													384
vaddubs	04				vD					vA					vB													512

**Table C-1. Instructions Sorted by Opcode in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vadduhs	04				vD					vA					vB														576
vadduws	04				vD					vA					vB														640
vaddsbs	04				vD					vA					vB														768
vaddshs	04				vD					vA					vB														832
vaddsws	04				vD					vA					vB														896
vsububm	04				vD					vA					vB														1024
vsubuhm	04				vD					vA					vB														1088
vsubuwm	04				vD					vA					vB														1152
vsubcuw	04				vD					vA					vB														1408
vsububs	04				vD					vA					vB														1536
vsubuhs	04				vD					vA					vB														1600
vsubuws	04				vD					vA					vB														1664
vsubsbs	04				vD					vA					vB														1792
vsubshs	04				vD					vA					vB														1856
vsubsws	04				vD					vA					vB														1920
vmaxub	04				vD					vA					vB														2
vmaxuh	04				vD					vA					vB														66
vmaxuw	04				vD					vA					vB														130
vmaxsb	04				vD					vA					vB														258
vmaxsh	04				vD					vA					vB														322
vmaxsw	04				vD					vA					vB														386
vminub	04				vD					vA					vB														514
vminuh	04				vD					vA					vB														578
vminuw	04				vD					vA					vB														642
vminsb	04				vD					vA					vB														770
vminsh	04				vD					vA					vB														834
vminsw	04				vD					vA					vB														898
vavgub	04				vD					vA					vB														1026
vavguh	04				vD					vA					vB														1090
vavguw	04				vD					vA					vB														1154
vavgusb	04				vD					vA					vB														1282
vavgsh	04				vD					vA					vB														1346
vavgsw	04				vD					vA					vB														1410

**Table C-1. Instructions Sorted by Opcode in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vrlb	04	vD			vA			vB			4																	
vrlh	04	vD			vA			vB			68																	
vrlw	04	vD			vA			vB			132																	
vs1b	04	vD			vA			vB			260																	
vs1h	04	vD			vA			vB			324																	
vs1w	04	vD			vA			vB			388																	
vsl	04	vD			vA			vB			452																	
vsrb	04	vD			vA			vB			516																	
vsrh	04	vD			vA			vB			580																	
vsrw	04	vD			vA			vB			644																	
vsr	04	vD			vA			vB			708																	
vsrab	04	vD			vA			vB			772																	
vsrah	04	vD			vA			vB			836																	
vsraw	04	vD			vA			vB			900																	
vand	04	vD			vA			vB			1028																	
vandc	04	vD			vA			vB			1092																	
vor	04	vD			vA			vB			1156																	
vxor	04	vD			vA			vB			1220																	
vnor	04	vD			vA			vB			1284																	
mfvscr	04	vD			0_0000			0000_0			1540																	
mtvscr	04	00_000			0_0000			vB			1604																	
vcmpqubx	04	vD			vA			vB			Rc	6																
vcmpquhx	04	vD			vA			vB			Rc	70																
vcmpquwx	04	vD			vA			vB			Rc	134																
vcmpqfpx	04	vD			vA			vB			Rc	198																
vcmpgefpx	04	vD			vA			vB			Rc	454																
vcmpgtubx	04	vD			vA			vB			Rc	518																
vcmpgtuhx	04	vD			vA			vB			Rc	582																
vcmpgtuwx	04	vD			vA			vB			Rc	646																
vcmpgtfpx	04	vD			vA			vB			Rc	710																
vcmpgtsbx	04	vD			vA			vB			Rc	774																
vcmpgtshx	04	vD			vA			vB			Rc	838																
vcmpgtswx	04	vD			vA			vB			Rc	902																

**Table C-1. Instructions Sorted by Opcode in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vcmpbfp	04				vD					vA					vB			Rc										966	
vmuloub	04				vD					vA					vB														8
vmulouh	04				vD					vA					vB														72
vmulosb	04				vD					vA					vB														264
vmulosh	04				vD					vA					vB														328
vmuleub	04				vD					vA					vB														520
vmuleuh	04				vD					vA					vB														584
vmulesb	04				vD					vA					vB														776
vmulesh	04				vD					vA					vB														840
vsum4ubs	04				vD					vA					vB														1544
vsum4sbs	04				vD					vA					vB														1800
vsum4shs	04				vD					vA					vB														1608
vsum2sws	04				vD					vA					vB														1672
vsumsws	04				vD					vA					vB														1928
vaddfp	04				vD					vA					vB														10
vsubfp	04				vD					vA					vB														74
vrefp	04				vD					0_0000					vB														266
vrsqrtefp	04				vD					0_0000					vB														330
vexpteft	04				vD					0_0000					vB														394
vlogeft	04				vD					0_0000					vB														458
vrfn	04				vD					0_0000					vB														522
vrfiz	04				vD					0_0000					vB														586
vrfip	04				vD					0_0000					vB														650
vrfim	04				vD					0_0000					vB														714
vcfux	04				vD					UIMM					vB														778
vcsfx	04				vD					UIMM					vB														842
vctuxs	04				vD					UIMM					vB														906
vctsxs	04				vD					UIMM					vB														970
vmaxfp	04				vD					vA					vB														1034
vminfp	04				vD					vA					vB														1098
vmrghb	04				vD					vA					vB														12
vmrghh	04				vD					vA					vB														76
vmrghw	04				vD					vA					vB														140



**Table C-1. Instructions Sorted by Opcode in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vmrglb	04	vD			vA			vB			268																	
vmrglh	04	vD			vA			vB			332																	
vmrglw	04	vD			vA			vB			396																	
vspltb	04	vD			UIMM			vB			524																	
vsplth	04	vD			UIMM			vB			588																	
vspltw	04	vD			UIMM			vB			652																	
vspltisb	04	vD			SIMM			0000_0			780																	
vspltish	04	vD			SIMM			0000_0			844																	
vspltisw	04	vD			SIMM			0000_0			908																	
vslo	04	vD			vA			vB			1036																	
vsro	04	vD			vA			vB			1100																	
vpkuhum	04	vD			vA			vB			14																	
vpkuwum	04	vD			vA			vB			78																	
vpkuhus	04	vD			vA			vB			142																	
vpkuwus	04	vD			vA			vB			206																	
vpkshus	04	vD			vA			vB			270																	
vpkswus	04	vD			vA			vB			334																	
vpkshss	04	vD			vA			vB			398																	
vpkswss	04	vD			vA			vB			462																	
vupkhsb	04	vD			0_0000			vB			526																	
vupkhsh	04	vD			0_0000			vB			590																	
vupklisb	04	vD			0_0000			vB			654																	
vupklish	04	vD			0_0000			vB			718																	
vpkpx	04	vD			vA			vB			782																	
vupkhp	04	vD			0_0000			vB			846																	
vupklpx	04	vD			0_0000			vB			974																	
lvsl	31	vD			A			B			6						0											
lvsr	31	vD			A			B			38						0											
dst	31	0	0_0	STRM			A			B			342						0									
dstt	31	1	0_0	STRM			A			B			342						0									
dstst	31	0	0_0	STRM			A			B			374						0									
dststt	31	1	0_0	STRM			A			B			374						0									
dss	31	0	0_0	STRM			0_0000			0000_0			822						0									

**Table C-1. Instructions Sorted by Opcode in Decimal Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>dssall</b>	31	1	0_0	STRM			0_0000			0000_0												822						0
<b>lvebx</b>	31			vD				A					B									71						0
<b>lvehx</b>	31			vD				A					B									39						0
<b>lviewx</b>	31			vD				A					B									0						0
<b>lvx</b>	31			vD				A					B									103						0
<b>lvxl</b>	31			vD				A					B									359						0
<b>stvebx</b>	31			vS				A					B									135						0
<b>stvehx</b>	31			vS				A					B									167						0
<b>stviewx</b>	31			vS				A					B									199						0
<b>stvx</b>	31			vS				A					B									231						0
<b>stvxl</b>	31			vS				A					B									487						0

# Appendix D Instructions Sorted by Opcode

## D.1 Instructions Sorted by Opcode in Binary Format

Table D-1 lists AltiVec instructions grouped by opcode in binary format.

Key:

 Reserved bits

**Table D-1. Instructions Sorted by Opcode in Binary Format**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vmhaddshs</b>	0001_00				vD					vA					vB					vC								10_0000
<b>vmhraddshs</b>	0001_00				vD					vA					vB					vC								10_0001
<b>vmladduhm</b>	0001_00				vD					vA					vB					vC								10_0010
<b>vmsumubm</b>	0001_00				vD					vA					vB					vC								10_0100
<b>vmsummbm</b>	0001_00				vD					vA					vB					vC								10_0101
<b>vmsumuhm</b>	0001_00				vD					vA					vB					vC								10_0110
<b>vmsumuhs</b>	0001_00				vD					vA					vB					vC								10_0111
<b>vmsumshm</b>	0001_00				vD					vA					vB					vC								10_1000
<b>vmsumshs</b>	0001_00				vD					vA					vB					vC								10_1001
<b>vsel</b>	0001_00				vD					vA					vB					vC								10_1010
<b>vperm</b>	0001_00				vD					vA					vB					vC								10_1011
<b>vsldoi</b>	0001_00				vD					vA					vB			0		SH								10_1100
<b>vmaddfp</b>	0001_00				vD					vA					vB													000_0010_1110
<b>vnmsubfp</b>	0001_00				vD					vA					vB					vC								10_1111
<b>vaddubm</b>	0001_00				vD					vA					vB													000_0000_0000
<b>vadduhm</b>	0001_00				vD					vA					vB													000_0100_0000
<b>vadduwm</b>	0001_00				vD					vA					vB													000_1000_0000
<b>vaddcuw</b>	0001_00				vD					vA					vB													001_1000_0000
<b>vaddubs</b>	0001_00				vD					vA					vB													010_0000_0000
<b>vadduhs</b>	0001_00				vD					vA					vB													010_0100_0000

**Table D-1. Instructions Sorted by Opcode in Binary Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vadduws	0001_00				vD					vA					vB							010_1000_0000							
vaddsbs	0001_00				vD					vA					vB							011_0000_0000							
vaddshs	0001_00				vD					vA					vB							011_0100_0000							
vaddsws	0001_00				vD					vA					vB							011_1000_0000							
vsububm	0001_00				vD					vA					vB							100_0000_0000							
vsubuhm	0001_00				vD					vA					vB							100_0100_0000							
vsubuwm	0001_00				vD					vA					vB							100_1000_0000							
vsubcuw	0001_00				vD					vA					vB							101_1000_0000							
vsububs	0001_00				vD					vA					vB							110_0000_0000							
vsubuhs	0001_00				vD					vA					vB							110_0100_0000							
vsubuws	0001_00				vD					vA					vB							110_1000_0000							
vsubsbs	0001_00				vD					vA					vB							111_0000_0000							
vsubshs	0001_00				vD					vA					vB							111_0100_0000							
vsubsws	0001_00				vD					vA					vB							111_1000_0000							
vmaxub	0001_00				vD					vA					vB							000_0000_0010							
vmaxuh	0001_00				vD					vA					vB							000_0100_0010							
vmaxuw	0001_00				vD					vA					vB							000_1000_0010							
vmaxsb	0001_00				vD					vA					vB							001_0000_0010							
vmaxsh	0001_00				vD					vA					vB							001_0100_0010							
vmaxsw	0001_00				vD					vA					vB							001_1000_0010							
vminub	0001_00				vD					vA					vB							010_0000_0010							
vminuh	0001_00				vD					vA					vB							010_0100_0010							
vminuw	0001_00				vD					vA					vB							010_1000_0010							
vminsb	0001_00				vD					vA					vB							011_0000_0010							
vminsh	0001_00				vD					vA					vB							011_0100_0010							
vminsw	0001_00				vD					vA					vB							011_1000_0010							
vavgub	0001_00				vD					vA					vB							100_0000_0010							
vavguh	0001_00				vD					vA					vB							100_0100_0010							
vavguw	0001_00				vD					vA					vB							100_1000_0010							
vavgsb	0001_00				vD					vA					vB							101_0000_0010							
vavgsh	0001_00				vD					vA					vB							101_0100_0010							
vavgsw	0001_00				vD					vA					vB							101_1000_0010							
vrlb	0001_00				vD					vA					vB							000_0000_0100							

**Table D-1. Instructions Sorted by Opcode in Binary Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
<b>vrh</b>	0001_00				vD					vA					vB							000_0100_0100							
<b>vrlw</b>	0001_00				vD					vA					vB							000_1000_0100							
<b>vslb</b>	0001_00				vD					vA					vB							001_0000_0100							
<b>vslh</b>	0001_00				vD					vA					vB							001_0100_0100							
<b>vslw</b>	0001_00				vD					vA					vB							001_1000_0100							
<b>vsl</b>	0001_00				vD					vA					vB							001_1100_0100							
<b>vsrb</b>	0001_00				vD					vA					vB							010_0000_0100							
<b>vsrh</b>	0001_00				vD					vA					vB							010_0100_0100							
<b>vsrw</b>	0001_00				vD					vA					vB							010_1000_0100							
<b>vsr</b>	0001_00				vD					vA					vB							010_1100_0100							
<b>vsrab</b>	0001_00				vD					vA					vB							011_0000_0100							
<b>vsrah</b>	0001_00				vD					vA					vB							011_0100_0100							
<b>vsraw</b>	0001_00				vD					vA					vB							011_1000_0100							
<b>vand</b>	0001_00				vD					vA					vB							100_0000_0100							
<b>vandc</b>	0001_00				vD					vA					vB							100_0100_0100							
<b>vor</b>	0001_00				vD					vA					vB							100_1000_0100							
<b>vxor</b>	0001_00				vD					vA					vB							100_1100_0100							
<b>vnor</b>	0001_00				vD					vA					vB							101_0000_0100							
<b>mfvscr</b>	0001_00				vD				0_0000					0000_0								110_0000_0100							
<b>mtvscr</b>	0001_00				00_000				0_0000					vB								110_0100_0100							
<b>vcmpequbx</b>	0001_00				vD					vA					vB			Rc				00_0000_0110							
<b>vcmpequhx</b>	0001_00				vD					vA					vB			Rc				00_0100_0110							
<b>vcmpequwx</b>	0001_00				vD					vA					vB			Rc				00_1000_0110							
<b>vcmpeqfpx</b>	0001_00				vD					vA					vB			Rc				00_1100_0110							
<b>vcmpgefpx</b>	0001_00				vD					vA					vB			Rc				01_1100_0110							
<b>vcmpgtubx</b>	0001_00				vD					vA					vB			Rc				10_0000_0110							
<b>vcmpgtuhx</b>	0001_00				vD					vA					vB			Rc				10_0100_0110							
<b>vcmpgtuwx</b>	0001_00				vD					vA					vB			Rc				10_1000_0110							
<b>vcmpgtfpx</b>	0001_00				vD					vA					vB			Rc				10_1100_0110							
<b>vcmpgtsbx</b>	0001_00				vD					vA					vB			Rc				11_0000_0110							
<b>vcmpgtshx</b>	0001_00				vD					vA					vB			Rc				11_0100_0110							
<b>vcmpgtswx</b>	0001_00				vD					vA					vB			Rc				11_1000_0110							
<b>vcmpbfpx</b>	0001_00				vD					vA					vB			Rc				11_1100_0110							

**Table D-1. Instructions Sorted by Opcode in Binary Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vmuloub	0001_00				vD					vA					vB						000_0000_1000								
vmulouh	0001_00				vD					vA					vB							000_0100_1000							
vmulosb	0001_00				vD					vA					vB							001_0000_1000							
vmulosh	0001_00				vD					vA					vB							001_0100_1000							
vmuleub	0001_00				vD					vA					vB							010_0000_1000							
vmuleuh	0001_00				vD					vA					vB							010_0100_1000							
vmulesb	0001_00				vD					vA					vB							011_0000_1000							
vmulesh	0001_00				vD					vA					vB							011_0100_1000							
vsum4ubs	0001_00				vD					vA					vB							110_0000_1000							
vsum4sbs	0001_00				vD					vA					vB							111_0000_1000							
vsum4shs	0001_00				vD					vA					vB							110_0100_1000							
vsum2sws	0001_00				vD					vA					vB							110_1000_1000							
vsumsws	0001_00				vD					vA					vB							111_1000_1000							
vaddfp	0001_00				vD					vA					vB							000_0000_1010							
vsubfp	0001_00				vD					vA					vB							000_0100_1010							
vrefp	0001_00				vD					0_0000					vB							001_0000_1010							
vrsqrtefp	0001_00				vD					0_0000					vB							001_0100_1010							
vexptefp	0001_00				vD					0_0000					vB							001_1000_1010							
vlogefp	0001_00				vD					0_0000					vB							001_1100_1010							
vrfin	0001_00				vD					0_0000					vB							010_0000_1010							
vrfiz	0001_00				vD					0_0000					vB							010_0100_1010							
vrfip	0001_00				vD					0_0000					vB							010_1000_1010							
vrfim	0001_00				vD					0_0000					vB							010_1100_1010							
vcfux	0001_00				vD					UIMM					vB							011_0000_1010							
vcsrfx	0001_00				vD					UIMM					vB							011_0100_1010							
vctuxs	0001_00				vD					UIMM					vB							011_1000_1010							
vctxsx	0001_00				vD					UIMM					vB							011_1100_1010							
vmaxfp	0001_00				vD					vA					vB							100_0000_1010							
vminfp	0001_00				vD					vA					vB							100_0100_1010							
vmrghb	0001_00				vD					vA					vB							000_0000_1100							
vmrghh	0001_00				vD					vA					vB							000_0100_1100							
vmrghw	0001_00				vD					vA					vB							000_1000_1100							
vmrglb	0001_00				vD					vA					vB							001_0000_1100							

### Table D-1. Instructions Sorted by Opcode in Binary Format (continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vmrglh	0001_00				vD					vA					vB						001_0100_1100								
vmrglw	0001_00				vD					vA					vB							001_1000_1100							
vspltb	0001_00				vD					UIMM					vB							010_0000_1100							
vsplth	0001_00				vD					UIMM					vB							010_0100_1100							
vspltw	0001_00				vD					UIMM					vB							010_1000_1100							
vspltisb	0001_00				vD					SIMM				0000_0								011_0000_1100							
vspltish	0001_00				vD					SIMM				0000_0								011_0100_1100							
vspltisw	0001_00				vD					SIMM				0000_0								011_1000_1100							
vslo	0001_00				vD					vA					vB							100_0000_1100							
vsro	0001_00				vD					vA					vB							100_0100_1100							
vpkuhum	0001_00				vD					vA					vB							000_0000_1110							
vpkuwum	0001_00				vD					vA					vB							000_0100_1110							
vpkuhus	0001_00				vD					vA					vB							000_1000_1110							
vpkuwus	0001_00				vD					vA					vB							000_1100_1110							
vpkshus	0001_00				vD					vA					vB							001_0000_1110							
vpkswus	0001_00				vD					vA					vB							001_0100_1110							
vpkshss	0001_00				vD					vA					vB							001_1000_1110							
vpkswss	0001_00				vD					vA					vB							001_1100_1110							
vupkhsb	0001_00				vD					0_0000					vB							010_0000_1110							
vupkhs	0001_00				vD					0_0000					vB							010_0100_1110							
vupklsb	0001_00				vD					0_0000					vB							010_1000_1110							
vupklsh	0001_00				vD					0_0000					vB							010_1100_1110							
vpkpx	0001_00				vD					vA					vB							0110000_1110							
vupkhp	0001_00				vD					0_0000					vB							011_0100_1110							
vupklp	0001_00				vD					0_0000					vB							011_1100_1110							
lvsl	0111_11				vD					A					B							000_0000_110							0
lvsr	0111_11				vD					A					B							000_0100_110							0
dst	0111_11	0	0_0	STRM						A					B							010_1010_110							0
dstt	0111_1	1	0_0	STRM						A					B							010_1010_110							0
dstst	0111_11	0	0_0	STRM						A					B							010_1110_110							0
dststt	0111_11	1	0_0	STRM						A					B							010_1110_110							0
dss	0111_11	0	0_0	STRM						0_0000					0000_0							110_0110_110							0
dssall	0111_11	1	0_0	STRM						0_0000					0000_0							110_0110_110							0

**Table D-1. Instructions Sorted by Opcode in Binary Format (continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>lvebx</b>	0111_11		vD			A			B			000_0000_111						0										
<b>lvehx</b>	0111_11		vD			A			B			000_0100_111						0										
<b>lvewx</b>	0111_11		vD			A			B			000_1000_111						0										
<b>lvx</b>	0111_11		vD			A			B			000_1100_111						0										
<b>lvxl</b>	0111_11		vD			A			B			010_1100_111						0										
<b>stvebx</b>	0111_11		vS			A			B			001_0000_111						0										
<b>stvehx</b>	0111_11		vS			A			B			001_0100_111						0										
<b>stvewx</b>	0111_11		vS			A			B			001_1000_111						0										
<b>stvx</b>	0111_11		vS			A			B			001_1100_111						0										
<b>stvxl</b>	0111_11		vS			A			B			011_1100_111						0										



# Appendix E Instructions Sorted by Form

## E.1 Instructions Sorted by Form

Table E-1 through Table E-4 list the AltiVec instructions grouped by form.

Key:

 Reserved bits

Table E-1. VA-Form

	OPCD	vD	vA	vB		vC	XO
	OPCD	vD	vA	vB	0	SH	XO

Specific Instructions																															
Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
vmhaddshs	04				vD					vA					vB						vC								32		
vmhraddshs	04				vD					vA					vB						vC								33		
vmladduhm	04				vD					vA					vB						vC								34		
vmsumubm	04				vD					vA					vB						vC								36		
vmsummbm	04				vD					vA					vB						vC								37		
vmsumuhm	04				vD					vA					vB						vC								38		
vmsumuhs	04				vD					vA					vB						vC								39		
vmsumshm	04				vD					vA					vB						vC								40		
vmsumshs	04				vD					vA					vB						vC								41		
vsel	04				vD					vA					vB						vC								42		
vperm	04				vD					vA					vB						vC								43		
vsldoi	04				vD					vA					vB			0			SH								44		
vmaddfp	04				vD					vA					vB						vC								46		
vnmsubfp	04				vD					vA					vB						vC								47		

**Table E-2. VX-Form**

	OPCD	vD	vA	vB	XO	
	OPCD	vD	0_0000	0000_0	XO	0
	OPCD	00_000	0_0000	vB	XO	0
	OPCD	vD	0_0000	vB	XO	
	OPCD	vD	UIMM	vB	XO	
	OPCD	vD	SIMM	0000_0	XO	

Specific Instructions																															
Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
vaddubm	04			vD					vA						vB															0	
vadduhm	04			vD					vA						vB															64	
vadduwm	04			vD					vA						vB															128	
vaddcuw	04			vD					vA						vB															384	
vaddubs	04			vD					vA						vB															512	
vadduhs	04			vD					vA						vB															576	
vadduws	04			vD					vA						vB															640	
vaddsbs	04			vD					vA						vB															768	
vaddshs	04			vD					vA						vB															832	
vaddsws	04			vD					vA						vB															896	
vsububm	04			vD					vA						vB															1024	
vsubuhm	04			vD					vA						vB															1088	
vsubuwm	04			vD					vA						vB															1152	
vsubcuw	04			vD					vA						vB															1408	
vsububs	04			vD					vA						vB															1536	
vsubuhs	04			vD					vA						vB															1600	
vsubuws	04			vD					vA						vB															1664	
vsubsbs	04			vD					vA						vB															1792	
vsubshs	04			vD					vA						vB															1856	
vsubsws	04			vD					vA						vB															1920	
vmaxub	04			vD					vA						vB															2	
vmaxuh	04			vD					vA						vB															66	
vmaxuw	04			vD					vA						vB															130	
vmaxsb	04			vD					vA						vB															258	
vmaxsh	04			vD					vA						vB															322	

Specific Instructions																															
Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
vmaxsw	04			vD				vA					vB																	386	
vminub	04			vD				vA					vB																	514	
vminuh	04			vD				vA					vB																	578	
vminuw	04			vD				vA					vB																	642	
vminsb	04			vD				vA					vB																	770	
vminsh	04			vD				vA					vB																	834	
vminsw	04			vD				vA					vB																	898	
vavgub	04			vD				vA					vB																	1026	
vavguh	04			vD				vA					vB																	1090	
vavguw	04			vD				vA					vB																	1154	
vavgusb	04			vD				vA					vB																	1282	
vavgsh	04			vD				vA					vB																	1346	
vavgsw	04			vD				vA					vB																	1410	
vrlb	04			vD				vA					vB																	4	
vrlh	04			vD				vA					vB																	68	
vrlw	04			vD				vA					vB																	132	
vslb	04			vD				vA					vB																	260	
vslh	04			vD				vA					vB																	324	
vslw	04			vD				vA					vB																	388	
vsl	04			vD				vA					vB																	452	
vsrb	04			vD				vA					vB																	516	
vsrh	04			vD				vA					vB																	580	
vsrw	04			vD				vA					vB																	644	
vsr	04			vD				vA					vB																	708	
vsrab	04			vD				vA					vB																	772	
vsrah	04			vD				vA					vB																	836	
vsraw	04			vD				vA					vB																	900	
vand	04			vD				vA					vB																	1028	
vandc	04			vD				vA					vB																	1092	
vor	04			vD				vA					vB																	1156	
vnor	04			vD				vA					vB																	1284	
mfvscr	04			vD				0_0000					0000_0																	1540	
mtvscr	04			00_000				0_0000					vB																	1604	

**Freescale Semiconductor, Inc.**  
Instructions Sorted by Form

Specific Instructions																															
Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
vmuloub	04			vD				vA					vB																8		
vmulouh	04			vD				vA					vB																72		
vmulosb	04			vD				vA					vB																264		
vmulosh	04			vD				vA					vB																328		
vmuleub	04			vD				vA					vB																520		
vmuleuh	04			vD				vA					vB																584		
vmulesb	04			vD				vA					vB																776		
vmulesh	04			vD				vA					vB																840		
vsum4ubs	04			vD				vA					vB																1544		
vsum4sbs	04			vD				vA					vB																1800		
vsum4shs	04			vD				vA					vB																1608		
vsum2sws	04			vD				vA					vB																1672		
vsumsws	04			vD				vA					vB																1928		
vaddfp	04			vD				vA					vB																10		
vsubfp	04			vD				vA					vB																74		
vrefp	04			vD				0_0000					vB																266		
vrsqrtefp	04			vD				0_0000					vB																330		
vexpteFP	04			vD				0_0000					vB																394		
vlogefp	04			vD				0_0000					vB																458		
vrfin	04			vD				0_0000					vB																522		
vrfiz	04			vD				0_0000					vB																586		
vrfip	04			vD				0_0000					vB																650		
vrfim	04			vD				0_0000					vB																714		
vcfux	04			vD				UIMM					vB																778		
vcfsx	04			vD				UIMM					vB																842		
vctuxs	04			vD				UIMM					vB																906		
vctxsx	04			vD				UIMM					vB																970		
vmaxfp	04			vD				vA					vB																1034		
vminfp	04			vD				vA					vB																1098		
vmrghb	04			vD				vA					vB																12		
vmrghh	04			vD				vA					vB																76		
vmrghw	04			vD				vA					vB																140		
vmrglb	04			vD				vA					vB																268		

Specific Instructions																															
Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
vmrglh	04			vD						vA					vB															332	
vmrglw	04			vD						vA					vB															396	
vspltb	04			vD						UIMM					vB														524		
vsplth	04			vD						UIMM					vB														588		
vspltw	04			vD						UIMM					vB														652		
vspltisb	04			vD						SIMM					0000_0														780		
vspltish	04			vD						SIMM					0000_0														844		
vspltisw	04			vD						SIMM					0000_0														908		
vslo	04			vD						vA					vB														1036		
vsro	04			vD						vA					vB														1100		
vpkuhum	04			vD						vA					vB														14		
vpkuwum	04			vD						vA					vB														78		
vpkuhus	04			vD						vA					vB														142		
vpkuwus	04			vD						vA					vB														206		
vpkshus	04			vD						vA					vB														270		
vpkswus	04			vD						vA					vB														334		
vpkshss	04			vD						vA					vB														398		
vpkswss	04			vD						vA					vB														462		
vupkhsb	04			vD						0_0000					vB														526		
vupkhsh	04			vD						0_0000					vB														590		
vupklsb	04			vD						0_0000					vB														654		
vupklsh	04			vD						0_0000					vB														718		
vpkpx	04			vD						vA					vB														782		
vupkhpX	04			vD						0_0000					vB														846		
vupklpx	04			vD						0_0000					vB														974		
vxor	04			vD						vA					vB														1220		

**Table E-3. X-Form**

	OPCD		vD		vA		vB		XO	0
	OPCD		vS		vA		vB		XO	0
	OPCD	T	0_0	STRM	A		B		XO	0

**Specific Instructions**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
dst	31	T	0_0	STRM	A	B																342						0
dstt	31	1	0_0	STRM	A	B																342						0
dstst	31	T	0_0	STRM	A	B																374						0
dststt	31	1	0_0	STRM	A	B																374						0
dss	31	A	0_0	STRM	0_0000	0000_0																822						0
dssall	31	1	0_0	STRM	0_0000	0000_0																822						0
lvebx	31			vD	A	B																7						0
lvehx	31			vD	A	B																39						0
lviewx	31			vD	A	B																71						0
lvsl	31			vD	A	B																6						0
lvsl	31			vD	A	B																38						0
lvx	31			vD	A	B																103						0
lvxl	31			vD	A	B																359						0
stvebx	31			vS	A	B																135						0
stvehx	31			vS	A	B																167						0
stviewx	31			vS	A	B																199						0
stvx	31			vS	A	B																231						0
stvxl	31			vS	A	B																487						0

**Table E-4. VXR-Form**

OPCD	vD	vA	vB	Rc	XO
------	----	----	----	----	----

**Specific Instructions**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vcmpbfp	04			vD	vA	vB																Rc							966
vcmpeqfp	04			vD	vA	vB																Rc							198
vcmpequb	04			vD	vA	vB																Rc							6
vcmpequh	04			vD	vA	vB																Rc							70
vcmpequw	04			vD	vA	vB																Rc							134
vcmpgef	04			vD	vA	vB																Rc							454
vcmpgtf	04			vD	vA	vB																Rc							710
vcmpgts	04			vD	vA	vB																Rc							774

## Specific Instructions

vcmpgtshx	04	vD	vA	vB	Rc	838
vcmpgtswx	04	vD	vA	vB	Rc	902
vcmpgtubx	04	vD	vA	vB	Rc	518
vcmpgtuhx	04	vD	vA	vB	Rc	582
vcmpgtuwx	04	vD	vA	vB	Rc	646





# Appendix F

## Instruction Set Legend

### F.1 Instruction Set Legend

Table F-1 provides general information on the AltiVec instruction set such as the architectural level, privilege level, and form.

**Table F-1. AltiVec Instruction Set Legend**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
dss		√				VX
dssall		√				VX
dst	√					VX
dstst		√				VX
dststt		√				VX
dstt		√				VX
lvebx	√					X
lvehx	√					X
lviewx	√					X
lvsl	√					X
lvslr	√					X
lvx	√					X
lvxl	√					X
mfvscr	√					VX
mtvscr	√					VX
stvebx	√					X
stvehx	√					X
stviewx	√					X
stvx	√					X
stvxl	√					X

**Table F-1. AltiVec Instruction Set Legend (continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
vaddcuw	√					VX
vaddfp	√					VX
vaddsbs	√					VX
vaddshs	√					VX
vaddsws	√					VX
vaddubm	√					VX
vaddubs	√					VX
vadduhm	√					VX
vadduhs	√					VX
vadduwm	√					VX
vadduws	√					VX
vand	√					VX
vandc	√					VX
vavgsb	√					VX
vavgsh	√					VX
vavgsw	√					VX
vavgub	√					VX
vavguh	√					VX
vavguw	√					VX
vcfux	√					VX
vcfsx	√					VX
vcmpbfp	√					VXR
vcmpeqfp	√					VXR
vcmpequb	√					VXR
vcmpequh	√					VXR
vcmpequw	√					VXR
vcmpgef	√					VXR
vcmpgtf	√					VXR
vcmpgtsb	√					VXR
vcmpgtsh	√					VXR
vcmpgtsw	√					VXR
vcmpgtub	√					VXR
vcmpgtuh	√					VXR

**Table F-1. AltiVec Instruction Set Legend (continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>vcmpgtuw</b>	√					VXR
<b>vctxs</b>	√					VX
<b>vctuxs</b>	√					VX
<b>vexptefp</b>	√					VX
<b>vlogefp</b>	√					VX
<b>vmaddfp</b>	√					VA
<b>vmaxfp</b>	√					VX
<b>vmaxsb</b>	√					VX
<b>vmaxsh</b>	√					VX
<b>vmaxsw</b>	√					VX
<b>vmaxub</b>	√					VX
<b>vmaxuh</b>	√					VX
<b>vmaxuw</b>	√					VX
<b>vmhaddshs</b>	√					VA
<b>vmhraddshs</b>	√					VA
<b>vminfp</b>	√					VX
<b>vminsb</b>	√					VX
<b>vminsh</b>	√					VX
<b>vminsw</b>	√					VX
<b>vminub</b>	√					VX
<b>vminuh</b>	√					VX
<b>vminuw</b>	√					VX
<b>vmladduhm</b>	√					VA
<b>vmrghb</b>	√					VX
<b>vmrghh</b>	√					VX
<b>vmrghw</b>	√					VX
<b>vmrglb</b>	√					VX
<b>vmrglh</b>	√					VX
<b>vmrglw</b>	√					VX
<b>vmsummbm</b>	√					VA
<b>vmsumshm</b>	√					VA
<b>vmsumshs</b>	√					VA
<b>vmsumubm</b>	√					VA

**Table F-1. AltiVec Instruction Set Legend (continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
vmsumuhm	√					VA
vmsumuhs	√					VA
vmulesb	√					VX
vmulesh	√					VX
vmuleub	√					VX
vmuleuh	√					VX
vmulosb	√					VX
vmulosh	√					VX
vmuloub	√					VX
vmulouh	√					VX
vnmsubfp	√					VA
vnor	√					VX
vor	√					VX
vperm	√					VA
vpkpx	√					VX
vpkshss	√					VX
vpkshus	√					VX
vpkswss	√					VX
vpkuhum	√					VX
vpkuhus	√					VX
vpkswus	√					VX
vpkuwum	√					VX
vpkuwus	√					VX
vrefp	√					VX
vrfim	√					VX
vrfin	√					VX
vrfip	√					VX
vrfiz	√					VX
vrlb	√					VX
vrlh	√					VX
vrlw	√					VX
vrsqrtefp	√					VX
vsel	√					VA

**Table F-1. AltiVec Instruction Set Legend (continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
vsl	√					VX
vslb	√					VX
vsldoi	√					VA
vslh	√					VX
vslo	√					VX
vslw	√					VX
vspltb	√					VX
vsplth	√					VX
vspltisb	√					VX
vspltish	√					VX
vspltisw	√					VX
vspltw	√					VX
vsr	√					VX
vsrab	√					VX
vsrah	√					VX
vsraw	√					VX
vsrb	√					VX
vsrh	√					VX
vsro	√					VX
vsrw	√					VX
vsubcuw	√					VX
vsubfp	√					VX
vsubsubs	√					VX
vsubshs	√					VX
vsubsws	√					VX
vsububm	√					VX
vsubuhm	√					VX
vsububs	√					VX
vsubuhs	√					VX
vsubuwm	√					VX
vsubuws	√					VX
vsumsws	√					VX
vsum2sws	√					VX

**Table F-1. AltiVec Instruction Set Legend (continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>vsum4sbs</b>	√					VX
<b>vsum4shs</b>	√					VX
<b>vsum4ubs</b>	√					VX
<b>vupkhpX</b>	√					VX
<b>vupkhsb</b>	√					VX
<b>vupkhsh</b>	√					VX
<b>vupkhpX</b>	√					VX
<b>vupklsh</b>	√					VX
<b>vupklpX</b>	√					VX
<b>vupklsb</b>	√					VX
<b>vupklsh</b>	√					VX
<b>vxor</b>	√					VX

# Appendix G

## User's Manual Revision History

This appendix provides a list of the major differences between the *AltiVec Programming Environments Manual*, Revision 0 and Revision 1. Note that the list only covers the major changes to the user's manual.

Only minor formatting upgrades comprised the changes in Revision 2.

No major changes were made top Revision 1.

The major changes to the *AltiVec Programming Environments Manual*, Revision 0, are as follows:

Section, Page	Change
2.1.2, Page 2-4	Replace Figure 2-4, "Saving/Restoring the AltiVec Context Register (VRSAVE)" with the following:

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Field	VR0	VR1	VR2	VR3	VR4	VR5	VR6	VR7	VR8	VR9	VR10	VR11	VR12	VR13	VR14	VR15
Reset	0000_0000_0000_0000															
R/W	R/W using <b>mf spr</b> or <b>mt spr</b> instructions															
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Field	VR16	VR17	VR18	VR19	VR20	VR21	VR22	VR23	VR24	VR25	VR26	VR27	VR28	VR29	VR30	VR31
Reset	0000_0000_0000_0000															
R/W	R/W using <b>mf spr</b> or <b>mt spr</b> instructions															
SPR	SPR256															

2.2, Page 2-9	Figure 2-10—The vector registers are 128 bits wide not 64 bits wide as shown.
---------------	-------------------------------------------------------------------------------

## Freescale Semiconductor, Inc. Changes

Section, Page No.

4.2.2.4, Page 4-20 Change Table 4-9 as follows:

- the mnemonic for Vector Round to Floating-Point Integer Nearest should be **vrfin** not fvrfin.
- the mnemonic for Vector Round to Floating-Point Integer toward Zero should be **vrfiz**, not fvrfiz.
- the mnemonic for Vector Round to Floating-Point Integer toward Positive Infinity should be **vrfip**, not fvrfip.
- the mnemonic for Vector Round to Floating-Point Integer toward Minus Infinity should be **vrfim**, not fvrfim.

6.2, Page 6-24 Change the **mfvscr** encoding as shown below (note: bit 31 is not 0):

04	vD	0 0 0 0 0	0 0 0 0 0	1540
0	5 6	10 11	15 16	20 21 31

6.2, Page 6-25 Change the **mtvscr** encoding as shown below (note: bit 31 is not 0):

04	0 0 0 0 0	0 0 0 0 0	vB	1604
0	5 6	10 11	15 16	20 21 31

A.1, Page A-2 Change the **mfvscr** encoding as shown below (note: bit 31 is not 0):

<b>mfvscr</b>	04	vD	0 0 0 0 0	0 0 0 0 0	1540
---------------	----	----	-----------	-----------	------

A.1, Page A-2 Change the **mtvscr** encoding as shown below (note: bit 31 is not 0 and vD should be vB):

<b>mtvscr</b>	04	0 0 0 0 0	0 0 0 0 0	vB	1604
---------------	----	-----------	-----------	----	------

A.2, Page A-9 Change the **mfvscr** encoding as shown below (note: bit 31 is not 0):

<b>mfvscr</b>	000100	vD	0 0 0 0 0	0 0 0 0 0	110 0000 0100
---------------	--------	----	-----------	-----------	---------------

A.2, Page A-9 Change the **mtvscr** encoding as shown below (note: bit 31 is not 0):

<b>mtvscr</b>	000100	0 0 0 0 0	0 0 0 0 0	vB	110 0100 0100
---------------	--------	-----------	-----------	----	---------------

A.3, Page A-14 Change the **mfvscr** encoding as shown below (note: bit 31 is not 0):

<b>mfvscr</b>	04	vD	0 0 0 0 0	0 0 0 0 0	1540
---------------	----	----	-----------	-----------	------

A.3, Page A-14 Change the **mtvscr** encoding as shown below (note: bit 31 is not 0):

<b>mtvscr</b>	04	0 0 0 0 0	0 0 0 0 0	vB	1604
---------------	----	-----------	-----------	----	------



## Glossary of Terms and Abbreviations

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this book. Some of the terms and definitions included in the glossary are reprinted from IEEE Std 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*, copyright ©1985 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE.

---

**A** **Architecture.** A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible *implementations*.

**Asynchronous exception.** *Exceptions* that are caused by events external to the processor's execution. In this document, the term 'asynchronous exception' is used interchangeably with the word *interrupt*.

**Atomic access.** A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The PowerPC architecture implements atomic accesses through the **lwarx/stwcx** instruction pair.

---

**B** **BAT (block address translation) mechanism.** A software-controlled array that stores the available block address translations on-chip.

**Beat.** A single state on the 603e bus interface that may extend across multiple bus cycles. A 603e transaction can be composed of multiple address or data *beats*.

**Biased exponent.** An *exponent* whose range of values is shifted by a constant (bias). Typically a bias is provided to allow a range of positive values to express a range that includes both positive and negative values.

**Big-endian.** A byte-ordering method in memory where the address *n* of a word corresponds to the *most-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the *most-significant byte*. See *Little-endian*.

**Block.** An area of memory that ranges from 128 Kbyte to 256 Mbyte whose size, translation, and protection attributes are controlled by the BAT mechanism.

**Boundedly undefined.** A characteristic of certain operation results that are not rigidly prescribed by the PowerPC architecture. Boundedly-undefined results for a given operation may vary among implementations and between execution attempts in the same implementation.

Although the architecture does not prescribe the exact behavior for when results are allowed to be boundedly undefined, the results of executing instructions in contexts where results are allowed to be *boundedly undefined* are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.

**Branch folding.** The replacement with target instructions of a branch instruction and any instructions along the not-taken path when a branch is either taken or predicted as taken.

**Branch prediction.** The process of guessing whether a branch will be taken. Such predictions can be correct or incorrect; the term ‘predicted’ as it is used here does not imply that the prediction is correct (successful). The PowerPC architecture defines a means for static branch prediction as part of the instruction encoding.

**Branch resolution.** The determination of whether a branch is taken or not taken. A branch is said to be resolved when the processor can determine which instruction path to take. If the branch is resolved as predicted, the instructions following the predicted branch that may have been speculatively executed can complete. If the branch is not resolved as predicted, instructions on the mispredicted path, and any results of speculative execution, are purged from the pipeline and fetching continues from the nonpredicted path.

**Burst.** A multiple-beat data transfer whose total size is typically equal to a cache block.

**Bus clock.** Clock that causes the bus state transitions.

**Bus master.** The owner of the address or data bus; the device that initiates or requests the transaction.

## C

**Cache.** High-speed memory containing recently accessed data or instructions (subset of main memory).

**Cache block.** A small region of contiguous memory that is copied from memory into a *cache*. The size of a cache block may vary among processors; the maximum block size is one *page*. In PowerPC processors, *cache coherency* is maintained on a cache-block basis. Note that the term ‘cache block’ is often used interchangeably with ‘cache line’.

**Cache coherency.** An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor’s cache.

**Cache flush.** An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush (**dcbf**) instruction.

**Caching-inhibited.** A memory update policy in which the *cache* is bypassed and the load or store is performed to or from main memory.

**Cast out.** A *cache block* that must be written to memory when a cache miss causes a cache block to be replaced.

**Changed bit.** One of two *page history bits* found in each *page table entry* (PTE). The processor sets the changed bit if any store is performed into the *page*. See also *Page access history bits* and *Referenced bit*.

**Clean.** An operation that causes a cache block to be written to memory, if modified, and then left in a valid, unmodified state in the cache.

**Clear.** To cause a bit or bit field to register a value of zero. See also *Set*.

**Context synchronization.** An operation that ensures that all instructions in execution complete past the point where they can produce an *exception*, that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are *fetched* and executed in the new context. Context synchronization may result from executing specific instructions (such as **isync** or **rfi**) or when certain events occur (such as an exception).

**Copy-back operation.** A cache operation in which a cache line is copied back to memory to enforce cache coherency. Copy-back operations consist of snoop push-out operations and cache cast-out operations.

---

## D

**Denormalized number.** A nonzero floating-point number whose exponent has a reserved value, usually the format's minimum, and whose explicit or implicit leading significand bit is zero.

**Direct-mapped cache.** A cache in which each main memory address can appear in only one location within the cache, operates more quickly when the memory request is a cache hit.

**Direct-store segment access.** An access to an I/O address space. The 603 defines separate memory-mapped and I/O address spaces, or segments, distinguished by the corresponding segment register T bit in the address translation logic of the 603. If the T bit is cleared, the memory reference is a normal memory-mapped access and can use the virtual memory management hardware of the 603. If the T bit is set, the memory reference is a direct-store access.

**Double-word swap.** AltiVec processors implement a double-word swap when moving quad words between vector registers and memory. The double word swap performs an additional swap to keep vector registers and memory consistent in little-endian mode. Double-word swap is referred to as 'swizzling' in the AltiVec technology architecture specification. This feature is not supported by the PowerPC architecture.

---

## E

**Effective address (EA).** The 32-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a *physical memory* address.

**Exception.** A condition encountered by the processor that requires special, supervisor-level processing.

**Exception handler.** A software routine that executes when an exception is taken. Normally, the exception handler corrects the condition that caused the exception, or performs some other meaningful task (that may include aborting the program that caused the exception). The address for each exception handler is identified by an exception vector offset defined by the architecture and a prefix selected via the MSR.

**Extended opcode.** A secondary opcode field generally located in instruction bits 21–30, that further defines the instruction type. All PowerPC instructions are one word in length. The most significant 6 bits of the instruction are the *primary opcode*, identifying the type of instruction. *See also* Primary opcode.

**Exclusive state.** MEI state (E) in which only one caching device contains data that is also in system memory.

**Execution synchronization.** A mechanism by which all instructions in execution are architecturally complete before beginning execution (appearing to begin execution) of the next instruction. Similar to context synchronization but doesn't force the contents of the instruction buffers to be deleted and refetched.

**Exponent.** In the binary representation of a floating-point number, the exponent is the component that normally signifies the integer power to which the value two is raised in determining the value of the represented number. *See also* *Biased exponent*.

## F

**Feed-forwarding.** A 603e feature that reduces the number of clock cycles that an execution unit must wait to use a register. When the source register of the current instruction is the same as the destination register of the previous instruction, the result of the previous instruction is routed to the current instruction at the same time that it is written to the register file. With feed-forwarding, the destination bus is gated to the waiting execution unit over the appropriate source bus, saving the cycles which would be used for the write and read.

**Fetch.** Retrieving instructions from either the cache or main memory and placing them into the instruction queue.

**Floating-point register (FPR).** Any of the 32 registers in the floating-point register file. These registers provide the source operands and destination results for floating-point instructions. Load instructions move data from memory to FPRs and store instructions move data from FPRs to memory. The FPRs are 64 bits wide and store floating-point values in double-precision format

**Floating-point unit.** The functional unit in the 603e processor responsible for executing all floating-point instructions.

**Flush.** An operation that causes a cache block to be invalidated and the data, if modified, to be written to memory.

**Fraction.** In the binary representation of a floating-point number, the field of the *significand* that lies to the right of its implied binary point.

**Fully associative.** Addressing scheme where every cache location (every byte) can have any possible address.

---

**G**

**General-purpose register (GPR).** Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.

**Guarded.** The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.

---

**H**

**Harvard architecture.** An architectural model featuring separate caches and other memory management resources for instructions and data.

**Hashing.** An algorithm used in the *page table* search process.

---

**I**

**IEEE 754.** A standard written by the Institute of Electrical and Electronics Engineers that defines operations and representations of binary floating-point numbers.

**Illegal instructions.** A class of instructions that are not implemented for a particular PowerPC processor. These include instructions not defined by the PowerPC architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations instructions that are defined only for 32-bit implementations are considered to be illegal instructions.

**Implementation.** A particular processor that conforms to the PowerPC architecture, but may differ from other architecture-compliant implementations for example in design, feature set, and implementation of *optional* features. The PowerPC architecture has many different implementations.

**Implementation-dependent.** An aspect of a feature in a processor's design that is defined by a processor's design specifications rather than by the PowerPC architecture.

**Implementation-specific.** An aspect of a feature in a processor's design that is not required by the PowerPC architecture, but for which the PowerPC architecture may provide concessions to ensure that processors that implement the feature do so consistently.

**Imprecise exception.** A type of *synchronous exception* that is allowed not to adhere to the precise exception model (see *Precise exception*). The PowerPC architecture allows only floating-point exceptions to be handled imprecisely.

**Inexact.** Loss of accuracy in an arithmetic operation when the rounded result differs from the infinitely precise value with unbounded range.

**Instruction queue.** A holding place for instructions fetched from the current instruction stream.

**Integer unit.** The functional unit in the 603e responsible for executing all integer instructions.

**In-order.** An aspect of an operation that adheres to a sequential model. An operation is said to be performed in-order if, at the time that it is performed, it is known to be required by the sequential execution model. See *Out-of-order*.

**Instruction latency.** The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.

**Instruction parallelism.** A feature of PowerPC processors that allows instructions to be processed in parallel.

**Interrupt.** An external signal that causes the 603e to suspend current execution and take a predefined exception.

---

**K**

**Key bits.** A set of key bits referred to as Ks and Kp in each segment register and each BAT register. The key bits determine whether supervisor or user programs can access a *page* within that *segment* or *block*.

**Kill.** An operation that causes a *cache block* to be invalidated without writing any modified data to memory.

---

**L**

**Latency.** The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.

**L2 cache.** See *Secondary cache*.

**Least-significant bit (lsb).** The bit of least value in an address, register, field, data element, or instruction encoding.

**Least-significant byte (LSB).** The byte of least value in an address, register, data element, or instruction encoding.

**Little-endian.** A byte-ordering method in memory where the address  $n$  of a word corresponds to the *least-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the *most-significant byte*. See *Big-endian*.

**Loop unrolling.** Loop unrolling provides a way of increasing performance by allowing more instructions to be issued in a clock cycle. The compiler replicates the loop body to increase the number of instructions executed between a loop branch.

## M

**Mantissa.** The decimal part of logarithm.

**MEI (modified/exclusive/invalid).** *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that the PowerPC architecture does not specify the implementation of a MEI protocol to ensure cache coherency.

**MESI (modified/exclusive/shared/invalid).** *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that the PowerPC architecture does not specify the implementation of a MESI protocol to ensure cache coherency.

**Memory access ordering.** The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.

**Memory-mapped accesses.** Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.

**Memory coherency.** An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.

**Memory consistency.** Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).



**Memory management unit (MMU).** The functional unit that is capable of translating an *effective* (logical) *address* to a physical address, providing protection mechanisms, and defining caching methods.

**Microarchitecture.** The hardware details of a microprocessor's design. Such details are not defined by the PowerPC architecture.

**Mnemonic.** The abbreviated name of an instruction used for coding.

**Modified state.** MEI state (M) in which one, and only one, caching device has the valid data for that address. The data at this address in external memory is not valid.

**Most-significant bit (msb).** The highest-order bit in an address, registers, data element, or instruction encoding.

**Most-significant byte (MSB).** The highest-order byte in an address, registers, data element, or instruction encoding.

**Munging.** A modification performed on an *effective address* that allows it to appear to the processor that individual aligned scalars are stored as *little-endian* values, when in fact it is stored in *big-endian* order, but at different byte addresses within double words. Note that munging affects only the effective address and not the byte order. Note also that this term is not used by the PowerPC architecture.

**Multiprocessing.** The capability of software, especially operating systems, to support execution on more than one processor at the same time.

---

## N

**NaN.** An abbreviation for not a number; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs and quiet NaNs.

**No-op.** No-operation. A single-cycle operation that does not affect registers or generate bus activity.

**Normalization.** A process by which a floating-point value is manipulated such that it can be represented in the format for the appropriate precision (single- or double-precision). For a floating-point value to be representable in the single- or double-precision format, the leading implied bit must be a 1.

---

## O

**OEA (operating environment architecture).** The level of the architecture that describes PowerPC memory management model, supervisor-level registers, synchronization requirements, and the

exception model. It also defines the time-base feature from a supervisor-level perspective. Implementations that conform to the PowerPC OEA also conform to the PowerPC UISA and VEA.

**Optional.** A feature, such as an instruction, a register, or an exception, that is defined by the PowerPC architecture but not required to be implemented.

**Out-of-order.** An aspect of an operation that allows it to be performed ahead of one that may have preceded it in the sequential model, for example, speculative operations. An operation is said to be performed out-of-order if, at the time that it is performed, it is not known to be required by the sequential execution model. See *In-order*.

**Out-of-order execution.** A technique that allows instructions to be issued and completed in an order that differs from their sequence in the instruction stream.

**Overflow.** An condition that occurs during arithmetic operations when the result cannot be stored accurately in the destination register(s). For example, if two 32-bit numbers are multiplied, the result may not be representable in 32 bits. Since the 32-bit registers of the 603e cannot represent this sum, an overflow condition occurs.

## P

---

**Page.** A region in memory. The OEA defines a page as a 4-Kbyte area of memory, aligned on a 4-Kbyte boundary.

**Page access history bits.** The *changed* and *referenced* bits in the PTE keep track of the access history within the page. The referenced bit is set by the MMU whenever the page is accessed for a read or write operation. The changed bit is set when the page is stored into. See *Changed bit* and *Referenced bit*.

**Page fault.** A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a *page* not currently resident in *physical memory*. On PowerPC processors, a page fault exception condition occurs when a matching, valid *page table entry* (PTE[V] = 1) cannot be located.

**Page table.** A table in memory is comprised of *page table entries*, or PTEs. It is further organized into eight PTEs per PTEG (page table entry group). The number of PTEGs in the page table depends on the size of the page table (as specified in the SDR1 register).

**Page table entry (PTE).** Data structures containing information used to translate *effective address* to physical address on a 4-Kbyte page basis. A PTE consists of 8 bytes of information in a 32-bit processor and 16 bytes of information in a 64-bit processor.

**Park.** The act of allowing a bus master to maintain bus mastership without having to arbitrate.

**Persistent data stream.** A data stream is considered to be persistent when it is expected to be loaded from frequently.

**Physical memory.** The actual memory that can be accessed through the system's memory bus.

**Pipelining.** A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

**Precise exceptions.** A category of exception for which the pipeline can be stopped so instructions that preceded the faulting instruction can complete and subsequent instructions can be flushed and redispached after exception handling has completed. See *Imprecise exceptions*.

**Primary opcode.** The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction.

**Program order.** The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache

**Protection boundary.** A boundary between *protection domains*.

**Protection domain.** A protection domain is a segment, a virtual page, a BAT area, or a range of unmapped effective addresses. It is defined only when the appropriate relocate bit in the MSR (IR or DR) is 1.

## Q

**Quad word.** A group of 16 contiguous locations starting at an address divisible by 16.

**Quiesce.** To come to rest. The processor is said to quiesce when an exception is taken or a **sync** instruction is executed. The instruction stream is stopped at the decode stage and executing instructions are allowed to complete to create a controlled context for instructions that may be

affected by out-of-order, parallel execution. See *Context synchronization*.

**Quiet NaN.** A type of *NaN* that can propagate through most arithmetic operations without signaling exceptions. A quiet NaN is used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when invalid. See *Signaling NaN*.

## R

**rA.** The rA instruction field is used to specify a GPR to be used as a source or destination.

**rB.** The rB instruction field is used to specify a GPR to be used as a source.

**rD.** The rD instruction field is used to specify a GPR to be used as a destination.

**rS.** The rS instruction field is used to specify a GPR to be used as a source.

**Real address mode.** An MMU mode when no address translation is performed and the *effective address* specified is the same as the physical address. The processor's MMU is operating in real address mode if its ability to perform address translation has been disabled through the MSR registers IR and/or DR bits.

**Record bit.** Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.

**Referenced bit.** One of two *page history bits* found in each *page table entry* (PTE). The processor sets the *referenced bit* whenever the page is accessed for a read or write. See also *Page access history bits*.

**Register indirect addressing.** A form of addressing that specifies one GPR that contains the address for the load or store.

**Register indirect with immediate index addressing.** A form of addressing that specifies an immediate value to be added to the contents of a specified GPR to form the target address for the load or store.

**Register indirect with index addressing.** A form of addressing that specifies that the contents of two GPRs be added together to yield the target address for the load or store.

**Rename register.** Temporary buffers used by instructions that have finished execution but have not completed.

**Reservation.** The processor establishes a reservation on a *cache block* of memory space when it executes an **lwarx** instruction to read a memory semaphore into a GPR.

**Reservation station.** A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.

**RISC (reduced instruction set computing).** An *architecture* characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.

## S

**Scan interface.** The 603e test interface.

**Secondary cache.** A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.

**Set (*v*).** To write a nonzero value to a bit or bit field; the opposite of *clear*. The term ‘set’ may also be used to generally describe the updating of a bit or bit field.

**Set (*n*).** A subdivision of a *cache*. Cacheable data can be stored in a given location in one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose *cache block* corresponding to that address was used least recently. See *Set-associative*.

**Set-associative.** Aspect of cache organization in which the cache space is divided into sections, called *sets*. The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.

**Shadowing.** Shadowing allows a register to be updated by instructions that are executed out of order without destroying machine state information.

**Signaling NaN.** A type of *NaN* that generates an invalid operation program exception when it is specified as arithmetic operands. See *Quiet NaN*.

**Significand.** The component of a binary floating-point number that consists of an explicit or implicit leading bit to the left of its implied binary point and a fraction field to the right.

**SIMD.** Single instruction stream, multiple data streams. A vector instruction can operate on several data elements within a single instruction in a single functional unit. SIMD is a way to work with all the data at once (in parallel), which can make execution faster.

**Simplified mnemonics.** Assembler mnemonics that represent a more complex form of a common operation.

**Slave.** The device addressed by a master device. The slave is identified in the address tenure and is responsible for supplying or latching the requested data for the master during the data tenure.

**Snooping.** Monitoring addresses driven by a bus master to detect the need for coherency actions.

**Snoop push.** Response to a snooped transaction that hits a modified cache block. The cache block is written to memory and made available to the snooping device.

**Splat.** A splat instruction will take one element and replicates (splats) that value into a vector register. The purpose being to have all elements have the same value so they can be used as a constant to multiply other vector registers.

**Split-transaction.** A transaction with independent request and response tenures.

**Split-transaction bus.** A bus that allows address and data transactions from different processors to occur independently.

**Stage.** The term ‘stage’ is used in two different senses, depending on whether the pipeline is being discussed as a physical entity or a sequence of events. In the latter case, a stage is an element in the pipeline during which certain actions are performed, such as decoding the instruction, performing an arithmetic operation, or writing back the results. Typically, the latency of a stage is one processor clock cycle. Some events, such as dispatch, write-back, and completion, happen instantaneously and may be thought to occur at the end of a stage. An instruction can spend multiple cycles in one stage. An integer multiply, for example, takes multiple cycles in the execute stage. When this occurs, subsequent instructions may stall. An instruction may also occupy more than one stage

simultaneously, especially in the sense that a stage can be seen as a physical resource—for example, when instructions are dispatched they are assigned a place in the CQ at the same time they are passed to the execute stage. They can be said to occupy both the complete and execute stages in the same clock cycle.

**Stall.** An occurrence when an instruction cannot proceed to the next stage.

**Static branch prediction.** Mechanism by which software (for example, compilers) can hint to the machine hardware about the direction a branch is likely to take.

**Sticky bit.** A bit that when *set* must be cleared explicitly.

**Superscalar machine.** A machine that can issue multiple instructions concurrently from a conventional linear instruction stream.

**Supervisor mode.** The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.

**Synchronization.** A process to ensure that operations occur strictly *in order*. See *Context synchronization* and *Execution synchronization*.

**Synchronous exception.** An *exception* that is generated by the execution of a particular instruction or instruction sequence. There are two types of synchronous exceptions, *precise* and *imprecise*.

**System memory.** The physical memory available to a processor.

---

## T

**Tenure.** The period of bus mastership. For the 603e, there can be separate address bus tenures and data bus tenures. A tenure consists of three phases: arbitration, transfer, and termination.

**TLB (translation lookaside buffer).** A cache that holds recently-used *page table entries*.

**Throughput.** The measure of the number of instructions that are processed per clock cycle.

**Tiny.** A floating-point value that is too small to be represented for a particular precision format, including *denormalized* numbers; they do not include  $\pm 0$ .

**Transaction.** A complete exchange between two bus devices. A transaction is typically comprised of an address tenure and one or more data tenures, which may overlap or occur separately from the address tenure. A transaction may be minimally comprised of an address tenure only.

**Transfer termination.** Signal that refers to both signals that acknowledge the transfer of individual beats (of both single-beat transfer and individual beats of a burst transfer) and to signals that mark the end of the tenure.

**Transient stream.** A data stream is considered to be transient when it is likely to be referenced from infrequently.

---

## U

**UIA (user instruction set architecture).** The level of the architecture to which user-level software should conform. The UIA defines the base user-level instruction set, user-level registers, data types, floating-point memory conventions and exception model as seen by user programs, and the memory and programming models.

**Underflow.** A condition that occurs during arithmetic operations when the result cannot be represented accurately in the destination register. For example, underflow can happen if two floating-point fractions are multiplied and the result requires a smaller *exponent* and/or *mantissa* than the single-precision format can provide. In other words, the result is too small to be represented accurately.

**User mode.** The operating state of a processor used typically by application software. In user mode, software can access only certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.

---

## V

**vA.** The vA instruction field is used to specify a vector register to be used as a source or destination.

**vB.** The vB instruction field is used to specify a vector register to be used as a source.

**vC.** The vC instruction field is used to specify a vector register to be used as a source.

**vD.** The vD instruction field is used to specify a vector register to be used as a destination.



**vS.** The vS instruction field is used to specify a vector register to be used as a source.

**VEA (virtual environment architecture).** The level of the *architecture* that describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time-base facility from a user-level perspective. *Implementations* that conform to the PowerPC VEA also adhere to the UISA, but may not necessarily adhere to the OEA.

**Vector.** The spatial parallel processing of short, fixed-length one-dimensional matrices performed by an execution unit.

**Vector Register (VR).** Any of the 32 registers in the vector register file. Each vector register is 128 bits wide. These registers can provide the source operands and destination results for AltiVec instructions.

**Virtual address.** An intermediate address used in the translation of an *effective address* to a physical address.

**Virtual memory.** The address space created using the memory management facilities of the processor. Program access to virtual memory is possible only when it coincides with *physical memory*.

## W

**Way.** A location in the cache that holds a cache block, its tags and status bits.

**Weak ordering.** A memory access model that allows bus operations to be reordered dynamically, which improves overall performance and in particular reduces the effect of memory latency on instruction throughput.

**Word.** A 32-bit data element.

**Write-back.** A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is *cast out* to make room for newer data.

**Write-through.** A cache memory update policy in which all processor write cycles are written to both the cache and memory.



## Index

### A

- Acronyms and abbreviated terms, list, xxiv
- Address bus
  - address calculation, 4-26
  - address modes, 1-9
  - address translation for streams, 5-7
- Alignment
  - aligned scalars, LE mode, 3-4
  - effective address, 4-26
  - load and store, 4-26
  - load instruction support, 4-29
  - memory access and vector register, 3-6
  - misaligned accesses, 3-1
  - misaligned vectors, 3-7
  - partially executed instructions, 5-10
  - quad-word data alignment, 3-7
  - rules, 3-4
- AltiVec technology
  - address modes, 1-9
  - cache overview, 1-12
  - exception handling, 1-12
  - features list, 1-4
  - features not defined, 1-6
  - instruction set, 1-11, 6-9, A-1-F-6
  - instruction set architecture support, 1-5
  - interelement operations, 1-9
  - intraelement operations, 1-9
  - levels of the PowerPC architecture, 1-5
  - operations supported, 1-9
  - overview, 1-3
  - PowerPC architecture extension, 1-2
  - programming model, 1-6
  - register file structure, 2-4
  - register set, 1-6, 2-4, 2-8
  - SIMD-style extension, 1-3, 1-7
  - structural overview, 1-4
- Arithmetic instructions
  - floating-point, 4-19
  - integer, 4-1

### B

- Big-endian mode
  - accessing a misaligned quad word, 3-8

- byte ordering, 1-7, 3-3
  - concept, 3-3
  - mapping, quad word, 3-3
  - misaligned vector, 3-7
  - mixed-endian systems, 3-12
- Block count, 5-2
- Block size, 5-2
- Block stride, 5-2
- Byte ordering
  - aligned scalars, LE mode, 3-4
  - big-endian mode, default, 3-3
  - concept, 3-2
  - default, 1-7
  - LE bit in MSR, 3-3
  - least-significant byte (LSB), 3-3
  - little-endian mode description, 3-3
  - most-significant byte (MSB), 3-3
  - quad-word example, 3-3

### C

- Cache
  - cache management instructions, 4-42
  - data stream touch, 5-2
  - dss instruction, 5-5
  - dst instruction, 5-2
  - dstst instruction, 5-4
  - dstt instruction, 5-4
  - overview, 1-12, 5-1
  - prefetch, software-directed, 5-2
  - prioritizing cache block replacement, 5-9
  - stopping streams, 5-5
  - storing to streams, 5-4
  - transient streams, 5-4
- Cache management instructions, 4-42
- Classes of instructions, 4-2
- Compare instructions
  - floating-point, 4-22
  - integer, 4-13, 4-14
- Computation modes
  - PowerPC architecture support, 4-2
- Conventions, xxiii
  - classes of instructions, 4-2
  - computation modes, 4-2

- execution model, 4-2
- memory addressing, 4-3
- operand conventions, 3-1
- terminology, xxvii
- CR (condition register)
  - bit fields, 2-8
  - CR6 field, compare instructions, 2-8
  - move to/from CR instructions, 4-40

## D

- Data organization, memory, 3-1
- Data stream, 5-2
- Double-word swap, 3-6

## E

- Echo cancellation, 1-2
- Effective address calculation
  - EA modifications, 3-5
  - loads and stores, 4-26
  - overview, 4-3
- Estimate instructions, 4-24
- Exceptions
  - data address breakpoint, 5-10
  - DSI exception, 5-10
  - exception behavior of prefetch streams, 5-6
  - exception handling, 1-12
  - floating-point exceptions, 3-14
  - invalid operation exception, 3-16
  - log of zero exception, 3-16
  - NaN operand exception, 3-15
  - overflow exception, 3-17
  - overview, 5-1
  - precise exceptions, 5-12
  - priorities, 5-12
  - synchronous exceptions, 5-12
  - unavailable exception, 5-10
  - underflow exception, 3-17
  - zero divide exception, 3-16
- Exclusive OR (XOR), 3-4
- Execution model
  - conventions, 4-2
  - floating-point, 3-12
- Extended mnemonics, *see* Simplified mnemonics

## F

- Features list
  - AltiVec technology features, 1-4
  - features not defined, 1-6
- Floating-point model
  - arithmetic instructions, 4-19
  - compare instructions, 4-22
  - division function, 4-18

- estimate instructions, 4-24
- exceptions, 3-14
- execution model, 3-12
- infinities, 3-14
- instructions, overview, 4-17
- Java mode, 3-13
- modes, 3-13
- multiply-add instructions, 4-20
- NaNs, 3-17
- non-Java mode, 3-14
- rounding mode, 3-14
- rounding/conversion instructions, 4-21
- square root functions, 4-19
- Formatting instructions, 4-31

## H

- High-order byte numbering, 1-8

## I

- Instructions
  - cache management instructions, 4-42
  - classes of instructions, 4-2
  - computation modes, 4-2
  - control flow, 4-31
  - conventions, xxvii, 6-2
  - detailed descriptions, 6-9-6-177
  - floating-point
    - arithmetic, 4-19
    - compare, 4-22
    - computational instructions, 3-12
    - division function, 4-18
    - estimate instructions, 4-24
    - multiply-add, 4-20
    - noncomputational instructions, 3-12
    - overview, 4-17
    - rounding/conversion, 4-21
    - square root functions, 4-19
  - format, lists, E-1
  - formats, 6-1
  - formatting instructions, 4-31
  - general information, F-1, G-1
  - integer
    - arithmetic, 4-1, 4-4
    - compare, 4-13, 4-14
    - load, 4-27
    - logical, 4-1, 4-15
    - rotate/shift, 4-16
    - store, 4-30
  - listed by format, E-1
  - listed by mnemonic, 6-9-6-177, A-1
  - listed by opcode, C-1, D-1
  - load and store
    - address generation, integer, 4-26

- integer load, 4-27
- integer store, 4-30
- memory addressing, 4-3
- memory control instructions, 4-41
- merge instructions, 4-34
- mnemonics, lists, A-1
- notations, 6-2
- opcodes, lists, C-1, D-1
- overview, 1-11
- pack instructions, 4-31
- partially executed instructions, 5-10
- permutation instructions, 4-31
- permute instructions, 4-36
- PowerPC instructions, list, A-1, B-1
- processor control instructions, 4-39
- quick reference, F-1, G-1
- select instruction, 4-36
- shift instructions, 4-37
- splat instructions, 4-35
- syntax conventions, xxvii, 6-2
- unpack instructions, 4-33
- vector integer, *see* integer
- Integer instructions
  - arithmetic instructions, 4-1, 4-4
  - compare instructions, 4-13, 4-14
  - logical instructions, 4-1, 4-15
  - rotate/shift instructions, 4-16
  - store instructions, 4-30
- Integer load instructions, 4-27
- Interelement operations, 1-9
- Intraelement operations, 1-9
- Invalid operation exception, 3-16

## J

- Java mode, 3-13

## L

- Little-endian mode
  - accessing a misaligned quad word, 3-10
  - byte ordering, 3-3
  - description, 3-3
  - mapping, quad word, 3-4
  - misaligned vector, 3-7
  - mixed-endian systems, 3-12
  - swapping, 3-6
- Load/store
  - address generation, integer, 4-26
  - integer load instructions, 4-27
  - integer store instructions, 4-30
- Log of zero exception, 3-16
- Logical instructions, integer, 4-1, 4-15
- Low-order byte numbering, 1-8

## M

- Mathematical predicates, 4-23
- Memory addressing, 4-3
- Memory control instructions, 4-41
- Memory management unit (MMU)
  - memory bandwidth, 5-1
  - overview, 1-12, 5-1
  - prefetch
    - data stream touch, 5-2
    - dss instruction, 5-5
    - dst instruction, 5-2
    - dstst instruction, 5-4
    - dstt instruction, 5-4
    - exception behavior, 5-6
    - software-directed, 5-2
    - stopping streams, 5-5
    - storing to streams, 5-4
    - transient streams, 5-4
- Memory operands, 4-3
- Memory sharing, 5-1
- Memory, data organization, 3-1
- Merge instructions, 4-34
- Misalignment
  - accessing a quad word
    - big-endian mode, 3-8
    - little-endian mode, 3-10
  - misaligned accesses, 3-1
  - misaligned vectors, 3-7
- Mixed-endian systems, 3-12
- Modulo mode, 4-4
- Move to/from CR instructions, 4-40
- MSR (machine state register)
  - bit settings, 2-9
  - LE bit, 3-3
- Multiply-add instructions, 4-20
- Munging, description, 3-4

## N

- NaN (not a number)
  - conversion to integer, 3-18
  - floating-point NaNs, 3-17
  - operand exception, 3-15
  - precedence, 3-18
  - production, 3-18
- Non-Java mode, 3-14

## O

- OEA (operating environment architecture)
  - definition, xx
  - programming model, 2-2
- Operands
  - conventions, description, 1-7, 3-1

floating-point conventions, 1-8  
 memory operands, 4-3  
 Operating environment architecture, *see* OEA  
 Operations  
   interelement operations, 1-9  
   intraelement operations, 1-9  
 Overflow exception, 3-17

## P

Pack instructions, 4-31  
 Permutation instructions, 4-31  
 Permute instructions, 4-36  
 PowerPC architecture support  
   computation modes, 4-2  
   execution model, 4-2  
   features summary  
     defined features, 1-4  
     features not defined, 1-6  
   instruction list, A-1, B-1  
   levels of the PowerPC architecture, 1-5  
   operating environment architecture, xx  
   programming model, 1-6  
   registers affected by AltiVec technology, 2-8  
   user instruction set architecture, xix, 1-5  
   virtual environment architecture, xix, 1-5  
 Prefetch, software-directed, 5-2  
 Processor control instructions, 4-39

## Q

QNaN arithmetic, 3-18

## R

Record bit (Rc), 6-2  
 Registers  
   CR, 2-8  
   overview, 1-6, 2-1  
   PowerPC register set, 2-1, 2-8  
   register file, 2-4  
   SRR0/SRR1, 2-10  
   VRs, 2-4  
   VRSERVE, 2-6  
   VSCR, 2-4  
 Rotate instructions, 4-16  
 Rounding/conversion instructions, FP, 4-21

## S

Saturation detection, 4-4  
 Scalars  
   aligned, LE mode, 3-4  
   loads and stores, 3-11  
   misaligned loads and stores, 3-11

Segment registers  
   T bit, Glossary-4  
 Select instruction, 4-36  
 Shift instructions, 4-16, 4-37  
 SIMD-style extension, 1-3, 1-7  
 Simplified mnemonics, 4-40  
 SNaN arithmetic, 3-18  
 Splat instructions, 4-35  
 SRR0/SRR1 (status save/restore registers), 2-10  
 Streams  
   address translation, 5-7  
   definition, 5-3  
   implementation assumptions, 5-9  
   synchronization, 5-7  
   usage notes, 5-7  
 Stride, 5-2  
 Swizzle, *see* Double-word swap  
 Synchronization streams, 5-7

## T

Terminology conventions, xxvii  
 Transient streams, 5-4

## U

UISA (user instruction set architecture), xix, 1-5  
   programming model, 2-2  
 Underflow exception, 3-17  
 Unpack instructions, 4-33  
 User instruction set architecture, *see* UISA

## V

VEA (virtual environment architecture)  
   definition, xix, 1-5  
   programming model, 2-2  
   user-level cache control instructions, 4-41  
 Vector formatting instructions, 4-31  
 Vector integer compare instructions, *see* Integer compare instructions  
 Vector merge instructions, 4-34  
 Vector pack instructions, 4-31  
 Vector permutation instructions, 4-31  
 Vector permute instructions, 4-36  
 Vector select instruction, 4-36  
 Vector shift instructions, 4-37  
 Vector splat instructions, 4-35  
 Vector unpack instructions, 4-33  
 Virtual environment architecture, *see* VEA  
 VRs (vector registers)  
   memory access alignment and VR, 3-6  
   register file, 2-4  
 VRSERVE register, 2-6  
 VSCR (vector status and control register), 2-4

**X**

XOR (exclusive OR), 3-4

**Z**

Zero divide exception, 3-16





# Freescale Semiconductor, Inc.

Overview	1
AltiVec Register Set	2
Operand Conventions	3
Addressing Modes and Instruction Set Summary	4
Cache, Exceptions, and Memory Management	5
AltiVec Instructions	6
Appendix A: Instruction Set Mnemonics - Decimal	A
Appendix B: Instruction Set Mnemonics - Binary	B
Appendix C: Opcodes - Decimal	C
Appendix D: Opcodes - Binary	D
Appendix E: Forms	E
Appendix F: Legends	F
Appendix G: Revision History	G
Glossary of Terms and Abbreviations	GLO
Index	IND

# Freescale Semiconductor, Inc.

1

Overview

2

Altivec Register Set

3

Operand Conventions

4

Addressing Modes and Instruction Set Summary

5

Cache, Exceptions, and Memory Management

6

Altivec Instructions

A

Appendix A: Instruction Set Mnemonics - Decimal

B

Appendix B: Instruction Set Mnemonics - Binary

C

Appendix C: Opcodes - Decimal

D

Appendix D: Opcodes - Binary

E

Appendix E: Forms

F

Appendix F: Legends

G

Appendix G: Revision History

GLO

Glossary of Terms and Abbreviations

IND

Index