# Signal Processing Engine (SPE) Programming Environments Manual: 

A Supplement to the EREF

## How to Reach Us:

## Home Page:

www.freescale.com
Web Support:
http://www.freescale.com/support
USA/Europe or Locations Not Listed:
Freescale Semiconductor, Inc.
Technical Information Center, EL516
2100 East Elliot Road
Tempe, Arizona 85284
+1-800-521-6274 or
+1-480-768-2130
www.freescale.com/support
Europe, Middle East, and Africa:
Freescale Halbleiter Deutschland GmbH
Technical Information Center
Schatzbogen 7
81829 Muenchen, Germany
+44 1296380456 (English)
+46 852200080 (English)
+49 8992103559 (German)
+33 169354848 (French)
www.freescale.com/support

## Japan:

Freescale Semiconductor Japan Ltd.
Headquarters
ARCO Tower 15F
1-8-1, Shimo-Meguro, Meguro-ku
Tokyo 153-0064
Japan
0120191014 or
+81 354379125
support.japan @ freescale.com

## Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
Technical Information Center
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T., Hong Kong
+800 26668080
support.asia@freescale.com
For Literature Requests Only:
Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405

Denver, Colorado 80217
+1-800 441-2447 or
+1-303-675-2140
Fax: +1-303-675-2150
LDCForFreescaleSemiconductor @ hibbertgroup.com

Freescale ${ }^{T M}$ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. The Power Architecture and Power.org word marks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. The PowerPC name is a trademark of IBM Corp. and is used under license. IEEE 754 is a registered trademark of the Institute of Electrical and Electronics Engineers, Inc. (IEEE). This product is not endorsed or approved by the IEEE. All other product or service names are the property of their respective owners.
© Freescale Semiconductor, Inc., 2008. Printed in the United States of America. All rights reserved.

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor 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 which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor 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 Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor 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 Freescale Semiconductor was negligent regarding the design or manufacture of the part.

## Contents

Paragraph Number Title
Page ..... Number
About This Book
Chapter 1
Overview
1.1 Overview ..... 1-1
1.2 Register Model ..... 1-2
1.2.1 SPE Instructions ..... 1-3
1.2.1.1 Embedded Vector and Scalar Floating-Point Instructions ..... 1-6
1.3 SPE and Embedded Floating-Point Exceptions and Interrupts ..... 1-6
Chapter 2
SPE Register Model
2.1 Overview ..... 2-1
2.2 Register Model ..... 2-1
2.2.1 General-Purpose Registers (GPRs) ..... 2-3
2.2.2 Accumulator Register (ACC) ..... 2-4
2.2.3 Signal Processing Embedded Floating-Point Status and Control Register (SPEFSCR) ..... 2-5
2.2.3.1 Interrupt Vector Offset Registers (IVORs) ..... 2-8
2.2.3.2 Exception Bit in the Exception Syndrome Register (ESR) ..... 2-8
2.2.3.3 Condition Register (CR) ..... 2-8
2.2.3.4 SPE Available Bit in the Machine State Register (MSR) ..... 2-9
Chapter 3
SPE and Embedded Floating-Point Instruction Model
3.1 Overview ..... 3-1
3.2 SPE Instruction Set ..... 3-1
3.2.1 SPE Data Formats ..... 3-2
3.2.1.1 Integer Format ..... 3-2
3.2.1.2 Fractional Format ..... 3-2
3.2.2 Computational Operations ..... 3-2
3.2.2.1 Data Formats and Register Usage. ..... 3-4
3.2.2.1.1 Signed Fractions ..... 3-4
3.2.2.1.2 SPE Integer and Fractional Operations ..... 3-4
3.2.2.1.3 SPE Instructions ..... 3-5
3.2.3SPE Simplified Mnemonics3-11
3.3 Embedded Floating-Point Instruction Set ..... 3-11
3.3.1 Embedded Floating-Point Operations ..... 3-12

## Contents

Paragraph
Title
Page ..... Number
Number
3.3.1.1 Operational Modes ..... 3-12
3.3.1.2 Floating-Point Data Formats ..... 3-12
3.3.1.3 Overflow and Underflow ..... 3-13
3.3.1.4 IEEE Std $754^{\mathrm{TM}}$ Compliance ..... 3-14
3.3.1.5 Sticky Bit Handling for Exception Conditions ..... 3-15
3.3.1.6 Implementation Options Summary ..... 3-15
3.3.1.7 Saturation, Shift, and Bit Reverse Models ..... 3-15
3.3.1.7.1 Saturation ..... 3-16
3.3.1.7.2 Shift Left ..... 3-16
3.3.1.7.3 Bit Reverse ..... 3-16
3.3.2 Embedded Vector and Scalar Floating-Point Instructions ..... 3-16
3.3.3 Load/Store Instructions ..... 3-18
3.3.3.1 Floating-Point Conversion Models ..... 3-18
Chapter 4
SPE/Embedded Floating-Point Interrupt Model
4.1 Overview ..... 4-1
4.2 SPE Interrupts ..... 4-1
4.2.1 Interrupt-Related Registers ..... 4-1
4.2.2 Alignment Interrupt ..... 4-2
4.2.3 SPE/Embedded Floating-Point Unavailable Interrupt ..... 4-2
4.2.4 SPE Embedded Floating-Point Interrupts ..... 4-3
4.2.4.1 Embedded Floating-Point Data Interrupt ..... 4-3
4.2.4.2 Embedded Floating-Point Round Interrupt ..... 4-3
4.3 Interrupt Priorities ..... 4-4
4.4 Exception Conditions ..... 4-4
4.4.1 Floating-Point Exception Conditions ..... 4-5
4.4.1.1 Denormalized Values on Input. ..... 4-5
4.4.1.2 Embedded Floating-Point Overflow and Underflow ..... 4-5
4.4.1.3 Embedded Floating-Point Invalid Operation/Input Errors ..... 4-5
4.4.1.4 Embedded Floating-Point Round (Inexact) ..... 4-6
4.4.1.5 Embedded Floating-Point Divide by Zero ..... 4-6
4.4.1.6 Default Results ..... 4-6
Chapter 5 Instruction Set
$5.1 \quad$ Notation ..... 5-1
5.2 Instruction Fields ..... 5-2
5.3 Description of Instruction Operations ..... 5-2

## Contents

Paragraph
Title
Page Number Number
5.3.1 SPE Saturation and Bit-Reverse Models ..... 5-4
5.3.1.1 Saturation ..... 5-4
5.3.1.2 Bit Reverse ..... 5-5
5.3.2 Embedded Floating-Point Conversion Models ..... 5-5
5.3.2.1 Common Embedded Floating-Point Functions ..... 5-6
5.3.2.1.1 32-Bit NaN or Infinity Test ..... 5-6
5.3.2.1.2 Signal Floating-Point Error ..... 5-6
5.3.2.1.3 Round a 32-Bit Value ..... 5-6
5.3.2.1.4 Round a 64-Bit Value ..... 5-7
5.3.2.2 Convert from Single-Precision Floating-Point to Integer Word with Saturation ..... 5-7
5.3.2.3 Convert from Double-Precision Floating-Point to Integer Word with Saturation... ..... 5-9
5.3.2.4 Convert from Double-Precision Floating-Point to Integer Double Word with Saturation ..... 5-10
5.3.2.5 Convert to Single-Precision Floating-Point from Integer Word with Saturation ..... 5-11
5.3.2.6 Convert to Double-Precision Floating-Point from Integer Word with Saturation. 5 ..... 5-12
5.3.2.7 Convert to Double-Precision Floating-Point from Integer Double Word with Saturation ..... 5-13
5.3.3 Integer Saturation Models ..... 5-14
5.3.4 Embedded Floating-Point Results ..... 5-14
5.4 Instruction Set ..... 5-15
Appendix A
Embedded Floating-Point Results Summary
Appendix B
SPE and Embedded Floating-Point Opcode Listings
B. 1 Instructions (Binary) by Mnemonic ..... B-1
B. 2 Instructions (Decimal and Hexadecimal) by Opcode ..... B-9
B. 3 Instructions by Form ..... B-16

## Contents

## Paragraph

Number
Title
Page
Number

## Figures

Figure
Number ..... Title
Page ..... Number
1-1 SPE Register Model ..... 1-2
1-2 Two-Element Vector Operations ..... 1-3
2-1 SPE Register Model ..... 2-1
2-2 Integer, Fractional, and Floating-Point Data Formats and GPR Usage ..... 2-2
2-3 32- and 64-Bit Register Elements and Bit-Numbering Conventions ..... 2-3
2-4 General Purpose Registers (GPR0-GRP31) ..... 2-4
2-5 Accumulator (ACC) ..... 2-4
2-6 Signal Processing and Embedded Floating-Point Status and Control Register (SPEFSCR) ..... 2-5
3-1 Two-Element Vector Operations ..... 3-3
3-2 Integer and Fractional Operations ..... 3-5
3-3 Floating-Point Data Format ..... 3-12
4-1 SPE Interrupt-Related Registers ..... 4-1
5-1 Instruction Description ..... 5-15
5-2 Vector Absolute Value (evabs) ..... 5-68
5-3 Vector Add Immediate Word (evaddiw) ..... 5-69
0-1 Vector Add Signed, Modulo, Integer to Accumulator Word (evaddsmiaaw) ..... 5-70
5-4 Vector Add Signed, Saturate, Integer to Accumulator Word (evaddssiaaw) ..... 5-71
5-5 Vector Add Unsigned, Modulo, Integer to Accumulator Word (evaddumiaaw) ..... 5-72
5-6 Vector Add Unsigned, Saturate, Integer to Accumulator Word (evaddusiaaw) ..... 5-73
5-7 Vector Add Word (evaddw) ..... 5-74
5-8 Vector AND (evand) ..... 5-75
5-9 Vector AND with Complement (evandc) ..... 5-76
5-10 Vector Compare Equal (evempeq) ..... 5-77
5-11 Vector Compare Greater Than Signed (evempgts) ..... 5-78
5-12 Vector Compare Greater Than Unsigned (evempgtu) ..... 5-79
5-13 Vector Compare Less Than Signed (evemplts) ..... 5-80
5-14 Vector Compare Less Than Unsigned (evempltu) ..... 5-81
5-15 Vector Count Leading Signed Bits Word (eventlsw) ..... 5-82
5-16 Vector Count Leading Zeros Word (eventlzw) ..... 5-83
5-17 Vector Divide Word Signed (evdivws) ..... 5-84
5-18 Vector Divide Word Unsigned (evdivwu) ..... 5-85
5-19 Vector Equivalent (eveqv) ..... 5-86
5-20 Vector Extend Sign Byte (evextsb) ..... 5-87
5-21 Vector Extend Sign Half Word (evextsh) ..... 5-88
5-22 evldd Results in Big- and Little-Endian Modes ..... 5-112
5-23 evlddx Results in Big- and Little-Endian Modes ..... 5-113
5-24 evldh Results in Big- and Little-Endian Modes ..... 5-114
5-25 evldhx Results in Big- and Little-Endian Modes ..... 5-115
5-26 evldw Results in Big- and Little-Endian Modes ..... 5-116
5-27 evldwx Results in Big- and Little-Endian Modes ..... 5-117

# Figures 

Figure ..... Number
Title
Page
5-28 evlhhesplat Results in Big- and Little-Endian Modes ..... 5-118
5-29 evlhhesplatx Results in Big- and Little-Endian Modes ..... 5-119
5-30 evlhhossplat Results in Big- and Little-Endian Modes ..... 5-120
5-31 evlhhossplatx Results in Big- and Little-Endian Modes ..... 5-121
5-32 evlhhousplat Results in Big- and Little-Endian Modes ..... 5-122
5-33 evlhhousplatx Results in Big- and Little-Endian Modes ..... 5-123
5-34 evlwhe Results in Big- and Little-Endian Modes ..... 5-124
5-35 evlwhex Results in Big- and Little-Endian Modes ..... 5-125
5-36 evlwhos Results in Big- and Little-Endian Modes ..... 5-126
5-37 evlwhosx Results in Big- and Little-Endian Modes ..... 5-127
5-38 evlwhou Results in Big- and Little-Endian Modes ..... 5-128
5-39 evlwhoux Results in Big- and Little-Endian Modes ..... 5-129
5-40 evlwhsplat Results in Big- and Little-Endian Modes ..... 5-130
5-41 evlwhsplatx Results in Big- and Little-Endian Modes ..... 5-131
5-42 evlwwsplat Results in Big- and Little-Endian Modes ..... 5-132
5-43 evlwwsplatx Results in Big- and Little-Endian Modes ..... 5-133
5-44 High Order Element Merging (evmergehi) ..... 5-134
5-45 High Order Element Merging (evmergehilo) ..... 5-135
5-46 Low Order Element Merging (evmergelo) ..... 5-136
5-47 Low Order Element Merging (evmergelohi) ..... 5-137
5-48 evmhegsmfaa (Even Form) ..... 5-138
5-49 evmhegsmfan (Even Form) ..... 5-139
5-50 evmhegsmiaa (Even Form) ..... 5-140
5-51 evmhegsmian (Even Form) ..... 5-141
5-52 evmhegumiaa (Even Form) ..... 5-142
5-53 evmhegumian (Even Form) ..... 5-143
5-54 Even Multiply of Two Signed Modulo Fractional Elements (to Accumulator) (evmhesmf) ..... 5-144
5-55 Even Form of Vector Half-Word Multiply (evmhesmfaaw) ..... 5-145
5-56 Even Form of Vector Half-Word Multiply (evmhesmfanw) ..... 5-146
5-57 Even Form for Vector Multiply (to Accumulator) (evmhesmi) ..... 5-147
5-58 Even Form of Vector Half-Word Multiply (evmhesmiaaw) ..... 5-148
5-59 Even Form of Vector Half-Word Multiply (evmhesmianw) ..... 5-149
5-60 Even Multiply of Two Signed Saturate Fractional Elements (to Accumulator) (evmhessf) ..... 5-150
5-61 Even Form of Vector Half-Word Multiply (evmhessfaaw) ..... 5-151
5-62 Even Form of Vector Half-Word Multiply (evmhessfanw) ..... 5-152
5-63 Even Form of Vector Half-Word Multiply (evmhessiaaw) ..... 5-153
5-64 Even Form of Vector Half-Word Multiply (evmhessianw) ..... 5-154
5-65 Vector Multiply Half Words, Even, Unsigned, Modulo, Integer (to Accumulator) (evmheumi) ..... 5-155

## Figures

Figure
Number Title
Page
Number
5-66 Even Form of Vector Half-Word Multiply (evmheumiaaw) ..... 5-156
5-67 Even Form of Vector Half-Word Multiply (evmheumianw) ..... 5-157
5-68 Even Form of Vector Half-Word Multiply (evmheusiaaw) ..... 5-158
5-69 Even Form of Vector Half-Word Multiply (evmheusianw) ..... 5-159
5-70 evmhogsmfaa (Odd Form) ..... 5-160
5-71 evmhogsmfan (Odd Form) ..... 5-161
5-72 evmhogsmiaa (Odd Form) ..... 5-162
5-73 evmhogsmian (Odd Form) ..... 5-163
5-74 evmhogumiaa (Odd Form) ..... 5-164
5-75 evmhogumian (Odd Form) ..... 5-165
5-76 Vector Multiply Half Words, Odd, Signed, Modulo, Fractional (to Accumulator) (evmhosmf) ..... 5-166
5-77 Odd Form of Vector Half-Word Multiply (evmhosmfaaw) ..... 5-167
5-78 Odd Form of Vector Half-Word Multiply (evmhosmfanw) ..... 5-168
5-79 Vector Multiply Half Words, Odd, Signed, Modulo, Integer (to Accumulator) (evmhosmi) ..... 5-169
5-80 Odd Form of Vector Half-Word Multiply (evmhosmiaaw) ..... 5-170
5-81 Odd Form of Vector Half-Word Multiply (evmhosmianw) ..... 5-171
5-82 Vector Multiply Half Words, Odd, Signed, Saturate, Fractional (to Accumulator) (evmhossf) ..... 5-173
5-83 Odd Form of Vector Half-Word Multiply (evmhossfaaw) ..... 5-174
5-84 Odd Form of Vector Half-Word Multiply (evmhossfanw) ..... 5-175
5-85 Odd Form of Vector Half-Word Multiply (evmhossiaaw) ..... 5-176
5-86 Odd Form of Vector Half-Word Multiply (evmhossianw) ..... 5-177
5-87 Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer (to Accumulator) (evmhoumi) ..... 5-178
5-88 Odd Form of Vector Half-Word Multiply (evmhoumiaaw) ..... 5-179
5-89 Odd Form of Vector Half-Word Multiply (evmhoumianw) ..... 5-180
5-90 Odd Form of Vector Half-Word Multiply (evmhousiaaw) ..... 5-181
5-91 Odd Form of Vector Half-Word Multiply (evmhousianw) ..... 5-182
5-92 Initialize Accumulator (evmra) ..... 5-183
5-93 Vector Multiply Word High Signed, Modulo, Fractional (to Accumulator) (evmwhsmf) ..... 5-184
5-94 Vector Multiply Word High Signed, Modulo, Integer (to Accumulator) (evmwhsm) ..... 5-185
5-95 Vector Multiply Word High Signed, Saturate, Fractional (to Accumulator) (evmwhssf) ..... 5-187
5-96 Vector Multiply Word High Unsigned, Modulo, Integer (to Accumulator) (evmwhumi) ..... 5-188
5-97 Vector Multiply Word Low Signed, Modulo, Integer and Accumulate in Words (evmwlsmiaaw) ..... 5-189

## Figures

Figure
Number ..... Title
Page ..... Number
5-98 Vector Multiply Word Low Signed, Modulo, Integer and Accumulate Negative in Words (evmwlsmianw) ..... 5-190
5-99 Vector Multiply Word Low Signed, Saturate, Integer and Accumulate in Words (evmwlssiaaw) ..... 5-192
5-100 Vector Multiply Word Low Signed, Saturate, Integer and Accumulate Negative in Words (evmwlssianw ..... 5-193
5-101 Vector Multiply Word Low Unsigned, Modulo, Integer (evmwlumi) ..... 5-194
5-102 Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate in Words (evmwlumiaaw) ..... 5-195
5-103 Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate Negative in Words (evmwlumianw) ..... 5-196
5-104 Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate in Words (evmwlusiaaw) ..... 5-197
5-105 Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate Negative in Words (evmwlusianw) ..... 5-198
5-106 Vector Multiply Word Signed, Modulo, Fractional (to Accumulator) (evmwsmf) ..... 5-199
5-107 Vector Multiply Word Signed, Modulo, Fractional and Accumulate (evmwsmfaa) ..... 5-200
5-108 Vector Multiply Word Signed, Modulo, Fractional and Accumulate Negative (evmwsmfan) ..... 5-201
5-109 Vector Multiply Word Signed, Modulo, Integer (to Accumulator) (evmwsmi) ..... 5-202
5-110 Vector Multiply Word Signed, Modulo, Integer and Accumulate (evmwsmiaa) ..... 5-203
5-111 Vector Multiply Word Signed, Modulo, Integer and Accumulate Negative (evmwsmian) ..... 5-204
5-112 Vector Multiply Word Signed, Saturate, Fractional (to Accumulator) (evmwssf) ..... 5-205
5-113 Vector Multiply Word Signed, Saturate, Fractional, and Accumulate (evmwssfaa) ..... 5-206
5-114 Vector Multiply Word Signed, Saturate, Fractional and Accumulate Negative (evmwssfan) ..... 5-207
5-115 Vector Multiply Word Unsigned, Modulo, Integer (to Accumulator) (evmwumi) ..... 5-208
5-116 Vector Multiply Word Unsigned, Modulo, Integer and Accumulate (evmwumiaa) ..... 5-209
5-117 Vector Multiply Word Unsigned, Modulo, Integer and Accumulate Negative (evmwumian) ..... 5-210
5-118 Vector NAND (evnand) ..... 5-211
5-119 Vector Negate (evneg) ..... 5-212

## Figures

Figure
Title
Page
Number ..... Title
Number
5-120 Vector NOR (evnor) ..... 5-213
5-121 Vector OR (evor) ..... 5-214
5-122 Vector OR with Complement (evorc) ..... 5-215
5-123 Vector Rotate Left Word (evrlw) ..... 5-216
5-124 Vector Rotate Left Word Immediate (evrlwi) ..... 5-217
5-125 Vector Round Word (evrndw) ..... 5-218
5-126 Vector Select (evsel) ..... 5-219
5-127 Vector Shift Left Word (evslw) ..... 5-220
5-128 Vector Shift Left Word Immediate (evslwi) ..... 5-221
5-129 Vector Splat Fractional Immediate (evsplatfi) ..... 5-222
5-130 evsplati Sign Extend ..... 5-223
5-131 Vector Shift Right Word Immediate Signed (evsrwis) ..... 5-224
5-132 Vector Shift Right Word Immediate Unsigned (evsrwiu) ..... 5-225
5-133 Vector Shift Right Word Signed (evsrws) ..... 5-226
5-134 Vector Shift Right Word Unsigned (evsrwu) ..... 5-227
5-135 evstdd Results in Big- and Little-Endian Modes ..... 5-228
5-136 evstddx Results in Big- and Little-Endian Modes ..... 5-229
5-137 evstdh Results in Big- and Little-Endian Modes ..... 5-230
5-138 evstdhx Results in Big- and Little-Endian Modes ..... 5-231
5-139 evstdw Results in Big- and Little-Endian Modes ..... 5-232
5-140 evstdwx Results in Big- and Little-Endian Modes ..... 5-233
5-141 evstwhe Results in Big- and Little-Endian Modes ..... 5-234
5-142 evstwhex Results in Big- and Little-Endian Modes ..... 5-235
5-143 evstwho Results in Big- and Little-Endian Modes ..... 5-236
5-144 evstwhox Results in Big- and Little-Endian Modes ..... 5-237
5-145 evstwwe Results in Big- and Little-Endian Modes ..... 5-238
5-146 evstwwex Results in Big- and Little-Endian Modes ..... 5-239
5-147 evstwwo Results in Big- and Little-Endian Modes ..... 5-240
5-148 evstwwox Results in Big- and Little-Endian Modes ..... 5-241
5-149 Vector Subtract Signed, Modulo, Integer to Accumulator Word (evsubfsmiaaw) ..... 5-242
5-150 Vector Subtract Signed, Saturate, Integer to Accumulator Word (evsubfssiaaw) ..... 5-243
5-151 Vector Subtract Unsigned, Modulo, Integer to Accumulator Word (evsubfumiaaw) ..... 5-244
5-152 Vector Subtract Unsigned, Saturate, Integer to Accumulator Word (evsubfusiaaw) ..... 5-245
5-153 Vector Subtract from Word (evsubfw) ..... 5-246
5-154 Vector Subtract Immediate from Word (evsubifw) ..... 5-247
5-155 Vector XOR (evxor) ..... 5-248

Figures
Figure Number

## Title

## Tables

Table Number Title
Page
Number
1-1 SPE Vector Multiply Instruction Mnemonic Structure ..... 1-1
1-2 Mnemonic Extensions for Multiply Accumulate Instructions ..... 1-4
1-3 SPE Vector Multiply Instruction Mnemonic Structure ..... 1-5
1-4 Mnemonic Extensions for Multiply-Accumulate Instructions ..... 1-5
2-1 SPEFSCR Field Descriptions ..... 2-5
2-2 SPE Instructions that Use the CR ..... 2-8
2-3 Embedded Floating-Point Instructions that Use the CR ..... 2-8
3-1 Mnemonic Extensions for Multiply Accumulate Instructions ..... 3-3
3-2 SPE Vector Multiply Instruction Mnemonic Structure ..... 3-5
3-3 Mnemonic Extensions for Multiply-Accumulate Instructions ..... 3-5
3-4 SPE Instructions ..... 3-6
3-5 SPE Simplified Mnemonics ..... 3-11
3-6 Vector and Scalar Floating-Point Instructions ..... 3-17
4-1 SPE/SPE Embedded Floating-Point Interrupt and Exception Types ..... 4-1
5-1 Notation Conventions ..... 5-1
5-2 Instruction Field Descriptions ..... 5-2
5-3 RTL Notation ..... 5-2
5-4 Operator Precedence ..... 5-4
5-5 Conversion Models ..... 5-5
5-6 Data Samples and Sizes ..... 5-16
A-1 Embedded Floating-Point Results Summary—Add, Sub, Mul, Div ..... A-1
A-2 Embedded Floating-Point Results Summary-Single Convert from Double ..... A-5
A-3 Embedded Floating-Point Results Summary-Double Convert from Single ..... A-5
A-4 Embedded Floating-Point Results Summary-Convert to Unsigned ..... A-6
A-5 Embedded Floating-Point Results Summary-Convert to Signed ..... A-6
A-6 Results Summary-Convert from Unsigned ..... A-6
A-7 Embedded Floating-Point Results Summary-Convert from Signed ..... A-7
A-8 Embedded Floating-Point Results Summary-*abs, *nabs, *neg ..... A-7
B-1 Instructions (Binary) by Mnemonic ..... B-1
B-2 Instructions (Decimal and Hexadecimal) by Opcode ..... B-9
B-3 Instructions (Binary) by Form ..... B-16

## Tables

## Table

Number
Title
Page
Number

## About This Book

The primary objective of this manual is to help programmers provide software compatible with processors that implement the signal processing engine (SPE) and embedded floating-point instruction sets.
To locate any published errata or updates for this document, refer to the web at http://www.freescale.com.
This book is used as a reference guide for assembler programmers. It 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.

The SPE Programming Interface Manual (SPEPIM) is a reference guide for high-level programmers. The VLEPIM describes how programmers can access SPE functionality from programming languages such as C and C++. It defines a programming model for use with the SPE instruction set. Processors that implement the Power ISA ${ }^{\text {TM }}$ (instruction set architecture) use the SPE instruction set as an extension to the base and embedded categories of the Power ISA.

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

- 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.
- Virtual environment architecture (VEA)—The VEA, which is the smallest component of the 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.
- 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.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Most of the discussions on the SPE are at the UISA level. For ease in reference, this book and the processor reference manuals have arranged the architecture information into topics that build on 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 devices that implement SPE. It must be kept in mind that each processor that implements the Power ISA 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.freescale.com.

## Audience

This manual is intended for system software and hardware developers, and for application programmers who want to develop products using the SPE. It is assumed that the reader understands operating systems, microprocessor system design, the basic principles of RISC processing, and details of the Power ISA.
This book describes how SPE interacts with the other components of the 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 SPE. This chapter provides an overview of how the VLE defines the register set, operand conventions, addressing modes, instruction set, and interrupt model.
- Chapter 2, "SPE Register Model," lists the register resources defined by the SPE and embedded floating-point ISAs. It also lists base category resources that are accessed by SPE and embedded floating-point instructions.
- Chapter 3, "SPE and Embedded Floating-Point Instruction Model," describes the SPE and embedded floating-point instruction set, including operand conventions, addressing modes, and instruction syntax. It also provides a brief description of instructions grouped by category.
- Chapter 5, "Instruction Set," functions as a handbook for the SPE and embedded floating-point 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.
- Appendix A, "Embedded Floating-Point Results Summary," summarizes the results of various types of embedded floating-point operations on various combinations of input operands.
- Appendix B, "SPE and Embedded Floating-Point Opcode Listings," lists all SPE and embedded-floating point instructions, grouped according to mnemonic and opcode.

This manual also includes 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 VLE and the Power ISA.

## General Information

The following documentation provides useful information about the Power Architecture ${ }^{\mathrm{TM}}$ technology and computer architecture in general:

- Computer Architecture: A Quantitative Approach, Third Edition, by John L. Hennessy and David A. Patterson.
- Computer Organization and Design: The Hardware/Software Interface, Third Edition, David A. Patterson and John L. Hennessy.


## Related Documentation

Freescale documentation is available from the sources listed on the back of the title page; the document order numbers, when applicable, are included in parentheses for ease in ordering:

- EREF: A Programmer's Reference Manual for Freescale Embedded Processors (EREFRM). Describes the programming, memory management, cache, and interrupt models defined by the Power ISA for embedded environment processors.
- Power ISA ${ }^{\mathrm{TM}}$. The latest version of the Power ISA can be downloaded from the website www.power.org.
- Variable-Length Encoding (VLE) Extension Programming Interface Manual (VLEPIM). Provides the VLE-specific extensions to the e500 application binary interface.
- e500 Application Binary Interface User's Guide (E500ABIUG). Establishes a standard binary interface for application programs on systems that implement the interfaces defined in the System V Interface Definition, Issue 3. This includes systems that have implemented UNIX System V Release 4.
- Reference manuals. The following reference manuals provide details information about processor cores and integrated devices:
- Core reference manuals-These books describe the features and behavior of individual microprocessor cores and provide specific information about how functionality described in the EREF is implemented by a particular core. They also describe implementation-specific features and microarchitectural details, such as instruction timing and cache hardware details, that lie outside the architecture specification.
- Integrated device reference manuals-These manuals describe the features and behavior of integrated devices that implement a Power ISA processor core. It is important to understand that some features defined for a core may not be supported on all devices that implement that core.
Also, some features are defined in a general way at the core level and have meaning only in the context of how the core is implemented. For example, any implementation-specific behavior of register fields can be described only in the reference manual for the integrated device.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Each of these documents include the following two chapters that are pertinent to the core:

- A core overview. This chapter provides a general overview of how the core works and indicates which of a core's features are implemented on the integrated device.
- A register summary chapter. This chapter gives the most specific information about how register fields can be interpreted in the context of the implementation.
These reference manuals also describe how the core interacts with other blocks on the integrated device, especially regarding topics such as reset, interrupt controllers, memory and cache management, debug, and global utilities.
- Addenda/errata to reference manuals-Errata documents are provided to address errors in published documents.
Because some processors have follow-on parts, often an addendum is provided that describes the additional features and functionality changes. These addenda, which may also contain errata, are intended for use with the corresponding reference manuals.
Always check the Freescale website for updates to reference manuals.
- Hardware specifications-Hardware specifications provide specific data regarding bus timing; signal behavior; AC, DC, and thermal characteristics; and other design considerations.
- Product brief-Each integrated device has a product brief that provides an overview of its features. This document is roughly the equivalent to the overview (Chapter 1) of the device's reference manual.
- Application notes-These short documents address specific design issues useful to programmers and engineers working with Freescale processors.

Additional literature is published as new processors become available. For current documentation, refer to http://www.freescale.com.

## 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 <br> one, it is said to be set. |
| :--- | :--- |
| mnemonics | Instruction mnemonics are shown in lowercase bold <br> italics |
|  | Italics indicate variable command parameters, for example, bcctr $x$ <br> Book titles in text are set in italics |
| $0 x 0$ | Prefix to denote hexadecimal number |
| 0 b0 | Prefix to denote binary number |
| rA, rB | Instruction syntax used to identify a source general-purpose register (GPR) |
| rD | Instruction syntax used to identify a destination GPR |
| frA, frB, frC | Instruction syntax used to identify a source floating-point register (FPR) |
| frD | Instruction syntax used to identify a destination FPR |
| REG[FIELD] | Abbreviations for registers are shown in uppercase text. Specific bits, fields, or <br> ranges appear in brackets. |

In some contexts, such as signal encodings, an unitalicized $x$ indicates a don't care.
$x$
$n$
ᄀ
\&
|

An italicized $x$ indicates an alphanumeric variable
An italicized $n$ indicates a numeric variable
NOT logical operator
AND logical operator
OR logical operator
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 5.1, "Notation."

## Acronyms and Abbreviations

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

Table i. Acronyms and Abbreviated Terms

| Term | Meaning |
| :---: | :--- |
| CR | Condition register |
| CTR | Count register |
| DEC | Decrementer register |
| EA | Effective address |
| EREF | A Programmer's Reference Manual for Freescale Embedded Processors (Including the e200 and e500 <br> Families) |
| GPR | General-purpose register |
| IEEE | Institute of Electrical and Electronics Engineers |
| IU | Integer unit |
| LR | Link register |
| LRU | Least recently used |
| LSB | Least significant byte |
| Isb | Least significant bit |
| LSU | Load/store unit |
| MMU | Memory management unit |
| MSB | Most significant byte |
| msb | Most significant bit |
| MSR | Machine state register |
| NaN | Not a number |
| No-op | No operation |
| OEA | Operating environment architecture |
|  |  |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table i. Acronyms and Abbreviated Terms (continued)

| Term | Meaning |
| :---: | :--- |
| PMCn | Performance monitor counter register |
| PVR | Processor version register |
| RISC | Reduced instruction set computing |
| RTL | Register transfer language |
| SIMM | Signed immediate value |
| SPR | Special-purpose register |
| SRR0 | Machine status save/restore register 0 |
| SRR1 | Machine status save/restore register 1 |
| 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 |
| VA | Virtual address |
| VEA | Virtual environment architecture |
| VLEPEM | Variable-Length Encoding (VLE) Programming Environments Manual |
| VLEPIM | Variable-Length Encoding (VLE) Extension Programming Interface Manual (VLEPIM) |
| 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 |  |
| :--- | :--- |
| Extended mnemonics | Simplified mnemonics |
| Fixed-point unit (FXU) | Integer unit (IU) |
| 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 |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

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

| The Architecture Specification | Equivalent to: |
| :--- | :--- |
| BA, BB, BT | crbA, crbB, crbD (respectively) |
| BF, BFA | crfD, crfS (respectively) |
| D | d |
| DS | ds |
| I, /I, I/I | $0 \ldots 0$ (shaded) |
| RA, RB, RT, RS | rA, rB, rD, rS (respectively) |
| SI | SIMM |
| U | IMM |
| UI | UIMM |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Chapter 1 Overview

This chapter provides a general description of the signal processing engine (SPE) and the SPE embedded floating-point resources defined as part of the Power ISA ${ }^{\text {TM }}$ (instruction set architecture).

### 1.1 Overview

The SPE is a 64-bit, two-element, single-instruction multiple-data (SIMD) ISA, originally designed to accelerate signal processing applications normally suited to DSP operation. The two-element vectors fit within GPRs extended to 64 bits. SPE also defines an accumulator register (ACC) to allow for back-to-back operations without loop unrolling. Like the VEC category, SPE is primarily an extension of Book I but identifies some resources for interrupt handling in Book III-E.

In addition to add and subtract to accumulator operations, the SPE supports a number of forms of multiply and multiply-accumulate operations, as well as negative accumulate forms. These instructions are summarized in Table 1-3. The SPE supports signed, unsigned, and fractional forms. For these instructions, the fractional form does not apply to unsigned forms, because integer and fractional forms are identical for unsigned operands.

Mnemonics for SPE instructions generally begin with the letters 'ev' (embedded vector).
Table 1-1. SPE Vector Multiply Instruction Mnemonic Structure

| Prefix | Multiply Element |  | Data Type Element |  | Accumulate Element |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| evm | ho he <br> hog <br> heg <br> wh <br> wl <br> whg <br> wig <br> w | half odd ( $16 \times 16->32$ ) <br> half even ( $16 \times 16->32$ ) <br> half odd guarded (16x16->32) <br> half even guarded ( $16 \times 16->32$ ) <br> word high (32x32->32) <br> word low (32x32->32) <br> word high guarded (32x32->32) <br> word low guarded (32x32->32) <br> word (32x32->64) | usi <br> umi <br> ssi <br> $\mathbf{s s f}^{1}$ <br> smi <br> smf ${ }^{1}$ | unsigned saturate integer unsigned modulo integer signed saturate integer signed saturate fractional signed modulo integer signed modulo fractional | a <br> aa <br> an <br> aaw <br> anw | write to ACC <br> write to ACC \& added ACC <br> write to ACC \& negate ACC <br> write to ACC \& ACC in words <br> write to ACC \& negate ACC in words |

${ }^{1}$ Low word versions of signed saturate and signed modulo fractional instructions are not supported. Attempting to execute an opcode corresponding to these instructions causes boundedly undefined results.

The SPE is part of the Power ISA specification (where it is referred to as the category SPE). Closely associated with the SPE are the embedded floating-point categories, which may be implemented if the SPE is implemented and which consist of the following:

- Single-precision scalar (SP.FS)
- Single-precision vector (SP.FV)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

- Double-precision scalar (SP.FD)

The embedded floating-point categories provide floating-point operations compatible with IEEE Std $754^{\mathrm{TM}}$ to power- and space-sensitive embedded applications. As is true for all SPE categories, rather than implementing separate register floating-point registers (FPRs), these categories share the GPRs used for integer operations, extending them to 64 bits to support the vector single-precision and scalar double-precision categories. These extended GPRs are described in Section 2.2.1, "General-Purpose Registers (GPRs)."

### 1.2 Register Model

Figure 1-1 shows the register resources defined by the Power ISA for the SPE and embedded floating-point operations. Note that SPE operations may also affect other registers defined by the Power ISA.


Figure 1-1. SPE Register Model
These registers are briefly described as follows:

- General-purpose registers (GPRs). Note especially that the SPE does not define a new register file but uses an extended version of the general-purpose registers (GPRs) implemented on all Power ISA devices. The GPRs are used as follows:
- SPE (not including the embedded floating-point instructions) treat the 64-bit GPRs as a two-element vector for 32-bit fractional and integer computation.
- Embedded scalar single-precision floating-point instructions use only the lower word of the GPRs for single-precision computation.
- Embedded vector single-precision instructions treat the 64-bit GPRs as a two-element vector for 32-bit single-precision computation.
- Embedded scalar double-precision floating-point instructions treat the GPRs as 64-bit single-element registers for double-precision computation.
- Accumulator register (ACC). Holds the results of the multiply accumulate (MAC) forms of SPE integer instructions. The ACC allows back-to-back execution of dependent MAC instructions, something that is found in the inner loops of DSP code such as finite impulse response (FIR) filters. The accumulator is partially visible to the programmer in that its results do not have to be explicitly read to use them. Instead, they are always copied into a 64-bit destination GPR specified as part of the instruction.
- SPE floating-point status and control register (SPEFSCR). Used for status and control of SPE and embedded floating-point instructions. It controls the handling of floating-point exceptions and records status resulting from the floating-point operations.
- Interrupt vector offset registers (IVORs). The SPE uses four IVORs, which together with the interrupt vector prefix register (IVPR) define the vector address for interrupt handler routines. The following IVORs are used:
- IVOR5 (SPR 405)—Defined by the base architecture for alignment exceptions and used with SPE load and store instructions alignment interrupts.
- IVOR32 (SPR 528)—SPE/embedded floating-point unavailable exception (causes the SPE/embedded floating-point unavailable interrupt)
- IVOR33 (SPR 529)—Embedded floating-point data interrupts
- IVOR34 (SPR 530)—Embedded floating-point round interrupts
- SPE/embedded floating-point available bit in the machine state register (MSR[SPV], formerly called MSR[SPE]). If this bit is zero and software attempts to execute an SPE/embedded floating-point instruction, an SPE unavailable interrupt is taken.
- Exception bit in the exception syndrome register (ESR[SPV], formerly called ESR[SPE). This bit is set whenever the processor takes an interrupt related to the execution of SPE vector or floating-point instructions.

Chapter 2, "SPE Register Model," provides detailed descriptions of these register resources.

### 1.2.1 SPE Instructions

.Instructions are provided for the instruction types:

- Simple vector instructions. These instructions use the corresponding low- and high-word elements of the operands to produce a vector result that is placed in the destination register, the accumulator, or both. Figure 1-2 shows how operations are typically performed in vector operations.


Figure 1-2. Two-Element Vector Operations

- Multiply and accumulate instructions. These instructions perform multiply operations, optionally add the result to the ACC, and place the result into the destination register and optionally into the ACC. These instructions are composed of different multiply forms, data formats, and data accumulate options, as indicated by their mnemonics, as shown in Table 1-2.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 1-2. Mnemonic Extensions for Multiply Accumulate Instructions

| Extension | Meaning | Comments |
| :---: | :---: | :---: |
| Multiply Form |  |  |
| he | Half word even | $16 \times 16 \rightarrow 32$ |
| heg | Half word even guarded | $16 \times 16 \rightarrow 32$, 64-bit final accum result |
| ho | Half word odd | $16 \times 16 \rightarrow 32$ |
| hog | Half word odd guarded | $16 \times 16 \rightarrow 32$, 64-bit final accum result |
| w | Word | $32 \times 32 \rightarrow 64$ |
| wh | Word high | $32 \times 32 \rightarrow 32$ (high order 32 bits of product) |
| wl | Word low | $32 \times 32 \rightarrow 32$ (low order 32 bits of product) |
| Data Format |  |  |
| smf | Signed modulo fractional | Modulo, no saturation or overflow |
| smi | Signed modulo integer | Modulo, no saturation or overflow |
| ssf | Signed saturate fractional | Saturation on product and accumulate |
| ssi | Signed saturate integer | Saturation on product and accumulate |
| umi | Unsigned modulo integer | Modulo, no saturation or overflow |
| usi | Unsigned saturate integer | Saturation on product and accumulate |
| Accumulate Option |  |  |
| a | Place in accumulator | Result $\rightarrow$ accumulator |
| aa | Add to accumulator | Accumulator + result $\rightarrow$ accumulator |
| aaw | Add to accumulator | Accumulator $_{0: 31}+$ result $_{0: 31} \rightarrow$ accumulator $_{0: 31}$ <br> Accumulator $_{32: 63}+$ result $_{32: 63} \rightarrow$ accumulator $_{32: 63}$ |
| an | Add negated to accumulator | Accumulator - result $\rightarrow$ accumulator |
| anw | Add negated to accumulator | Accumulator $_{0: 31}$ - result $_{0: 31} \rightarrow$ accumulator ${ }_{0: 31}$ Accumulator $_{32: 63}$ - result $_{32: 63} \rightarrow$ accumulator $_{32: 63}$ |

- Load and store instructions. These instructions provide load and store capabilities for moving data to and from memory. A variety of forms are provided that position data for efficient computation.
- Compare and miscellaneous instructions. These instructions perform miscellaneous functions such as field manipulation, bit reversed incrementing, and vector compares.

SPE supports several different computational capabilities. Modulo results produce truncation of the overflow bits in a calculation; therefore, overflow does not occur and no saturation is performed. For instructions for which overflow occurs, saturation provides a maximum or minimum representable value (for the data type) in the case of overflow.

Table 1-3 shows how SPE vector multiply instruction mnemonics are structured.
Table 1-3. SPE Vector Multiply Instruction Mnemonic Structure

| Prefix |  | Multiply Element | Data Type Element |  | Accumulate Element |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| evm | ho he hog heg wh wl whg wig w | half odd ( $16 \times 16->32$ ) <br> half even ( $16 \times 16->32$ ) <br> half odd guarded (16x16->32) <br> half even guarded ( $16 \times 16->32$ ) <br> word high (32×32->32) <br> word low (32x32->32) <br> word high guarded ( $32 \times 32->32$ ) <br> word low guarded (32x32->32) <br> word (32×32->64) | usi umi ssi ssf ${ }^{1}$ smi smf ${ }^{1}$ | unsigned saturate integer unsigned modulo integer signed saturate integer signed saturate fractional signed modulo integer signed modulo fractional | a <br> aa <br> an <br> aaw anw | write to ACC <br> write to ACC \& added ACC <br> write to ACC \& negate ACC <br> write to ACC \& ACC in words write to ACC \& negate ACC in words |

1 Low word versions of signed saturate and signed modulo fractional instructions are not supported. Attempting to execute an opcode corresponding to these instructions causes boundedly undefined results.

Table 1-4 defines mnemonic extensions for these instructions.
Table 1-4. Mnemonic Extensions for Multiply-Accumulate Instructions

| Extension | Meaning | Comments |
| :---: | :---: | :---: |
| Multiply Form |  |  |
| he | Half word even | $16 \times 16 \rightarrow 32$ |
| heg | Half word even guarded | $16 \times 16 \rightarrow 32,64$-bit final accumulator result |
| ho | Half word odd | $16 \times 16 \rightarrow 32$ |
| hog | Half word odd guarded | $16 \times 16 \rightarrow 32$, 64-bit final accumulator result |
| w | Word | $32 \times 32 \rightarrow 64$ |
| wh | Word high | $32 \times 32 \rightarrow 32$, high-order 32 bits of product |
| wl | Word Iow | $32 \times 32 \rightarrow 32$, low-order 32 bits of product |
| Data Type |  |  |
| smf | Signed modulo fractional | (Wrap, no saturate) |
| smi | Signed modulo integer | (Wrap, no saturate) |
| ssf | Signed saturate fractional |  |
| ssi | Signed saturate integer |  |
| umi | Unsigned modulo integer | (Wrap, no saturate) |
| usi | Unsigned saturate integer |  |
| Accumulate Options |  |  |
| a | Update accumulator | Update accumulator (no add) |
| aa | Add to accumulator | Add result to accumulator (64-bit sum) |
| aaw | Add to accumulator (words) | Add word results to accumulator words (pair of 32-bit sums) |
| an | Add negated | Add negated result to accumulator (64-bit sum) |
| anw | Add negated to accumulator (words) | Add negated word results to accumulator words (pair of 32-bit sums) |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 1.2.1.1 Embedded Vector and Scalar Floating-Point Instructions

The embedded floating-point operations are IEEE 754-compliant with software exception handlers and offer a simpler exception model than the Power ISA floating-point instructions that use the floating-point registers (FPRs). Instead of FPRs, these instructions use GPRs to offer improved performance for converting between floating-point, integer, and fractional values. Sharing GPRs allows vector floating-point instructions to use SPE load and store instructions.

Section 3.3.1.2, "Floating-Point Data Formats," describes the floating-point data format.

### 1.3 SPE and Embedded Floating-Point Exceptions and Interrupts

The SPE defines the following exceptions:

- SPE/embedded floating-point unavailable exception (causes the SPE/embedded floating-point unavailable interrupt)—IVOR32 (SPR 528)
- SPE vector alignment exception (causes the alignment interrupt)—IVOR5 (SPR 405)

In addition to these general SPE interrupts, the SPE embedded floating-point facility defines the following:

- Embedded floating-point data interrupt-IVOR33 (SPR 529)
- Embedded floating-point round interrupt—IVOR34 (SPR 539)

Details about these interrupts are provided in Chapter 4, "SPE/Embedded Floating-Point Interrupt Model."

## Chapter 2 <br> SPE Register Model

This chapter describes the register model of the signal processing engine (SPE) for embedded processors. This includes additional resources defined to support embedded floating-point instruction sets that may be implemented.

### 2.1 Overview

The SPE is designed to accelerate signal-processing applications normally suited to DSP operation. This is accomplished using short (two-element) vectors within 64-bit GPRs and using single instruction multiple data (SIMD) operations to perform the requisite computations. An accumulator register (ACC) allows back-to-back operations without loop unrolling.

### 2.2 Register Model

Figure 2-1 shows the register resources defined by the Power ISA for the SPE and embedded floating-point operations. Note that SPE operations may also affect other registers defined by the Power ISA.


Figure 2-1. SPE Register Model
Figure 2-2 shows how the SPE register model is used with the SPE and embedding floating-point instruction sets.


Note: Gray text indicates that this register or register field is not used.
1 Formatting of floating-point operands is as defined by IEEE 754.
Figure 2-2. Integer, Fractional, and Floating-Point Data Formats and GPR Usage

Several conventions regarding nomenclature are used in this chapter:

- All register bit numbering is 64-bit. As shown in Figure 2-3, for 64-bit registers, bit 0 being the most significant bit (msb). For 32-bit registers, bit 32 is the msb. For both 32-and 64-bit registers, bit 63 is the least significant bit (lsb).

64-bit register


Figure 2-3. 32- and 64-Bit Register Elements and Bit-Numbering Conventions

- As shown in Figure 2-3, bits 0 to 31 of a 64-bit register are referenced as the upper-, even-, or high-word element. Bits 32-63 are referred to as lower-, odd-, or low-word element.
- As shown in Figure 2-3, bits 0 to 15 and bits 32 to 47 are referenced as even half words. Bits 16 to 31 and bits 48 to 63 are odd half words.
- The gray lines shown in Figure 2-3 indicate 4-bit nibbles, and are provided as a convenience for making binary-to-hexadecimal conversions.
- Mnemonics for SPE instructions generally begin with the letters 'ev' (embedded vector).


### 2.2.1 General-Purpose Registers (GPRs)

The SPE requires a GPR file with thirty-two 64-bit registers, as shown in Figure 2-4, which also indicates how the SPE and embedded floating-point instruction sets use the GPRs. For 32-bit implementations, instructions that normally operate on a 32-bit register file access and change only the least significant 32 bits of the GPRs, leaving the most significant 32 bits unchanged. For 64-bit implementations, operation of these instructions is unchanged; that is, those instructions continue to operate on the 64-bit registers as they would if SPE were not implemented. SPE vector instructions view the 64-bit register as being composed of a vector of two 32-bit elements. (Some instructions read or write 16-bit elements.) The most significant 32 bits are called the upper, high, or even word. The least significant 32 bits are called the lower, low, or odd word. Unless otherwise specified, SPE instructions write all 64 bits of the destination register.


Figure 2-4. General Purpose Registers (GPR0-GRP31)
As shown in Figure 2-2 and Figure 2-4, embedded floating-point operations use the GPRs as follows:

- Single-precision floating-point requires a GPR file with thirty-two 32-bit or 64-bit registers. When implemented with a 64-bit register file on a 32-bit implementation, single-precision floating-point operations only use and modify bits $32-63$ of the GPR. In this case, bits $0-31$ of the GPR are left unchanged by a single-precision floating-point operation. For 64-bit implementations, bits $0-31$ are undefined after a single-precision floating-point operation.
- Vector floating-point and double-precision floating-point require a GPR file with thirty-two 64-bit GPRs.
- Floating-point double-precision instructions operate on the entire 64 bits of the GPRs where a floating-point data item consists of 64 bits.
- Vector floating-point instructions operate on the entire 64 bits of the GPRs, but contain two 32-bit data items that are operated on independently of each other in a SIMD fashion. The format of both data items is the same as a single-precision floating-point value. The data item contained in bits $0-31$ is called the "high word." The data item contained in bits $32-63$ is called the "low word."

There are no record forms of embedded floating-point instructions. Floating-point compare instructions treat NaNs, infinity, and denorm as normalized numbers for the comparison calculation when default results are provided.

### 2.2.2 Accumulator Register (ACC)

The 64-bit accumulator (ACC), shown in Figure 2-5, is used for integer/fractional multiply accumulate (MAC) forms of instructions. The ACC holds the results of the multiply accumulate forms of SPE fixed-point instructions. It allows the back-to-back execution of dependent MAC instructions, something that is found in the inner loops of DSP code such as FIR and FFT filters. It is partially visible to the programmer in that its results do not have to be explicitly read to be used. Instead they are always copied into a 64-bit destination GPR, specified as part of the instruction. Based on the instruction, the ACC can hold a single 64-bit value or a vector of two 32-bit elements.


Figure 2-5. Accumulator (ACC)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 2.2.3 Signal Processing Embedded Floating-Point Status and Control Register (SPEFSCR)

The SPEFSCR, shown in Figure 2-6, is used with SPE and embedded floating-point instructions. Vector floating-point instructions affect both the high element (bits 34-39) and low element floating-point status flags (bits 50-55). Double- and single-precision scalar floating-point instructions affect only the low-element floating-point status flags and leave the high-element floating-point status flags undefined.

SPR 512


Figure 2-6. Signal Processing and Embedded Floating-Point Status and Control Register (SPEFSCR)

Table 2-1 describes SPEFSCR bits.
Table 2-1. SPEFSCR Field Descriptions

| Bits | Name | Description |
| :---: | :---: | :--- |
| 32 | SOVH | Summary integer overflow high. Set when an SPE instruction sets OVH. This is a sticky bit that remains set <br> until it is cleared by an mtspr instruction. |
| 33 | OVH | Integer overflow high. OVH is set to indicate that an overflow occurred in the upper element during execution <br> of an SPE instruction. It is set if a result of an operation performed by the instruction cannot be represented in <br> the number of bits into which the result is to be placed and is cleared otherwise. OVH is not altered by modulo <br> instructions or by other instructions that cannot overflow. |
| 34 | FGH | Embedded floating-point guard bit high. Used by the floating-point round interrupt handler. FGH is an extension <br> of the low-order bits of the fractional result produced from a floating-point operation on the high word. FGH is <br> zeroed if an overflow, underflow, or invalid input error is detected on the high element of a vector floating-point <br> instruction. <br> Execution of a scalar floating-point instruction leaves FGH undefined. |
| 35 | FXH | Embedded floating-point inexact bit high. Used by the floating-point round interrupt handler. FXH is an <br> extension of the low-order bits of the fractional result produced from a floating-point operation on the high word. <br> FXH represents the logical OR of all of the bits shifted right from the guard bit when the fractional result is <br> normalized. FXH is zeroed if an overflow, underflow, or invalid input error is detected on the high element of a <br> vector floating-point instruction. <br> Execution of a scalar floating-point instruction leaves FXH undefined. |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 2-1. SPEFSCR Field Descriptions (continued)

| Bits | Name | $\quad$ Description |
| :---: | :---: | :--- |
| 36 | FINVH | Embedded floating-point invalid operation/input error high. Set under any of the following conditions: <br> - Any operand of a high word vector floating-point instruction is infinity, NaN, or denorm <br> - The operation is a divide and the dividend and divisor are both 0 <br> Execution of a scalar floating-point instruction leaves FINVH undefined. |
| 37 | FDBZH | Embedded floating-point divide by zero high. Set when a vector floating-point divide instruction is executed with <br> a divisor of 0 in the high word operand and the dividend is a finite non-zero number. <br> Execution of a scalar floating-point instruction leaves FDBZH undefined. |
| 38 | FUNFH | Embedded floating-point underflow high. Set when execution of a vector floating-point instruction results in an <br> underflow on the high word operation. <br> Execution of a scalar floating-point instruction leaves FUNFH undefined. |
| 39 | FOVFH | Embedded floating-point overflow high. Set when the execution of a vector floating-point instruction results in <br> an overflow on the high word operation. <br> Execution of a scalar floating-point instruction leaves FOVFH undefined. |
| $40-41$ | - | Reserved, should be cleared. |
| 42 | FINXS | Embedded floating-point inexact sticky flag. Set under the following conditions: <br> - Execution of any scalar or vector floating-point instruction delivers an inexact result for either the low or high <br> element ,and no floating-point data interrupt is taken for either element. |
| 49 | O A floating-point instruction results in overflow (FOVF=1 or FOVFH=1), but floating-point overflow exceptions |  |
| are disabled (FOVFE=0). |  |  |
| • A floating-point instruction results in underflow (FUNF=1 or FUNFH=1), but floating-point underflow |  |  |
| exceptions are disabled (FUNFE=0), and no floating-point data interrupt occurs. |  |  |
| FINXS is a sticky bit; it remains set until it is cleared by software. |  |  |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 2-1. SPEFSCR Field Descriptions (continued)

| Bits | Name | Description |
| :---: | :---: | :---: |
| 51 | FX | Embedded floating-point inexact bit (low/scalar). Used by the embedded floating-point round interrupt handler. FX is an extension of the low-order bits of the fractional result produced from an embedded floating-point instruction on the low word. FX represents the logical OR of all the bits shifted right from the guard bit when the fractional result is normalized. FX is zeroed if an overflow, underflow, or invalid input error is detected on embedded floating-point instruction. |
| 52 | FINV | Embedded floating-point invalid operation/input error (low/scalar). Set by one of the following: <br> - Any operand of a low-word vector or scalar floating-point operation is infinity, NaN , or denorm. <br> - The dividend and divisor are both 0 for a divide operation. <br> - A conversion to integer or fractional value overflows. |
| 53 | FDBZ | Embedded floating-point divide by zero (low/scalar). Set when an embedded floating-point divide instruction is executed with a divisor of 0 in the low word operand, and the dividend is a finite nonzero number. |
| 54 | FUNF | Embedded floating-point underflow (low/scalar). Set when the execution of an embedded floating-point instruction results in an underflow on the low word operation. |
| 55 | FOVF | Embedded floating-point overflow (Low/scalar). Set when the execution of an embedded floating-point instruction results in an overflow on the low word operation. |
| 56 | - | Reserved, should be cleared. |
| 57 | FINXE | Embedded floating-point round (inexact) exception enable <br> 0 Exception disabled <br> 1 Exception enabled. A floating-point round interrupt is taken if no other interrupt is taken, and if FG \| FGH | FX I FXH (signifying an inexact result) is set as a result of a floating-point operation. If a floating-point instruction operation results in overflow or underflow and the corresponding underflow or overflow exception is disabled, a floating-point round interrupt is taken. |
| 58 | FINVE | Embedded floating-point invalid operation/input error exception enable <br> 0 Exception disabled <br> 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FINV or FINVH. |
| 59 | FDBZE | Embedded floating-point divide by zero exception enable <br> 0 Exception disabled <br> 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FDBZ or FDBZH. |
| 60 | FUNFE | Embedded floating-point underflow exception enable <br> 0 Exception disabled <br> 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FUNF or FUNFH. |
| 61 | FOVFE | Embedded floating-point overflow exception enable <br> 0 Exception disabled <br> 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FOVF or FOVFH. |
| 62-63 | FRMC | Embedded floating-point rounding mode control <br> 00 Round to nearest <br> 01 Round toward zero <br> 10 Round toward +infinity. If this mode is not implemented, embedded floating-point round interrupts are generated for every floating-point instruction for which rounding is indicated. <br> 11 Round toward -infinity. If this mode is not implemented, embedded floating-point round interrupts are generated for every floating-point instruction for which rounding is indicated. |

${ }^{1}$ Software note: Software can detect the hardware that manages this bit by performing an operation on a NaN and observing whether hardware sets this sticky bit. Alternatively, if it desired that software work on all processors supporting embedded floating-point, software should check the appropriate status bits and set the sticky bit. If hardware also performs this operation, the action is redundant.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 2.2.3.1 Interrupt Vector Offset Registers (IVORs)

The SPE uses four IVORs which, together with the interrupt vector prefix register (IVPR), define the vector address for interrupt handler routines. The following IVORs are used:

- IVOR5 (SPR 405)—Defined by the base architecture for alignment interrupts and used for SPE load and store instructions alignment interrupts
- IVOR32 (SPR 528)—SPE/embedded floating-point unavailable exception (causes the SPE/embedded floating-point unavailable interrupt)
- IVOR33 (SPR 529)—Embedded floating-point data interrupts
- IVOR34 (SPR 530)—Embedded floating-point round interrupts

For more information, see Chapter 4, "SPE/Embedded Floating-Point Interrupt Model."

### 2.2.3.2 Exception Bit in the Exception Syndrome Register (ESR)

$\operatorname{ESR}[S P V]$ (ESR[56]), formerly called ESR[SPE], is set whenever the processor takes an interrupt related to the execution of SPE vector or floating-point instructions.

### 2.2.3.3 Condition Register (CR)

The CR is used to record results for compare and test instructions. It also provides a source operand for the Vector Select (evsel) instruction. Table 2-2 lists SPE instructions that explicitly access CR bits (crS or crD).

Table 2-2. SPE Instructions that Use the CR

| Instruction | Mnemonic | Syntax |
| :--- | :---: | :---: |
| Vector Compare Equal | evcmpeq | crD,rA,rB |
| Vector Compare Greater Than Signed | evcmpgts | crD,rA,rB |
| Vector Compare Greater Than Unsigned | evcmpgtu | crD,rA,rB |
| Vector Compare Less Than Signed | evcmplts | crD,rA,rB |
| Vector Compare Less Than Unsigned | evcmpltu | crD,rA,rB |
| Vector Select | evsel | rD,rA,rB,crS |

Table 2-2 lists embedded floating-point instructions that explicitly access CR bits (crD).
Table 2-3. Embedded Floating-Point Instructions that Use the CR

| Instruction | Single-Precision |  | Double- Precision Scalar | Syntax |
| :--- | :---: | :---: | :---: | :---: |
|  | Scalar | Vector |  |  |
| Floating-Point Compare Equal | efscmpeq | evfscmpeq | efdcmpeq | crD,rA,rB |
| Floating-Point Compare Greater Than | efscmpgt | evfscmpgt | efdcmpgt | crD,rA,rB |
| Floating-Point Compare Less Than | efscmplt | evfscmplt | efdcmplt | crD,rA,rB |
| Floating-Point Test Equal | efststeq | evfststeq | efdtsteq | crD,rA,rB |
| Floating-Point Test Greater Than | efststgt | evfststgt | efdtstgt | crD,rA,rB |
| Floating-Point Test Less Than | efststlt | evfststlt | efdtstlt | crD,rA,rB |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 2.2.3.4 SPE Available Bit in the Machine State Register (MSR)

MSR[SPV] (MSR[38]), formerly called MSR[SPE], is the SPE/embedded floating-point available bit. If this bit is zero and software attempts to execute an SPE instruction, an SPE unavailable interrupt is taken.

## NOTE (Software)

Software can use MSR[SPV] to detect when a process uses the upper 32 bits of a 64-bit register on a 32-bit implementation and thus save them on context switch.

## Chapter 3 <br> SPE and Embedded Floating-Point Instruction Model

This chapter describes the instruction model of the signal processing engine (SPE) for embedded processors. This includes additional resources defined to support embedded floating-point instruction sets that may be implemented.

Chapter 5, "Instruction Set," gives complete descriptions of individual SPE and embedded floating-point instructions. Section 5.3.1, "SPE Saturation and Bit-Reverse Models," provides pseudo-RTL for saturation and bit reversal to more accurately describe those functions that are referenced in the instruction pseudo-RTL.

### 3.1 Overview

The SPE is designed to accelerate signal-processing applications normally suited to DSP operation. This is accomplished using short (two-element) vectors within 64-bit GPRs and using single instruction multiple data (SIMD) operations to perform the requisite computations. An accumulator register (ACC) allows back-to-back operations without loop unrolling.

The SPE defines both computational and load store instructions. SPE load store instructions are necessary for 32-bit implementation to access 64-bit operands.

Embedded floating-point instructions, which may be implemented if the SPE is implemented, include the following computational instructions:

- Embedded vector single-precision floating-point, which use extended 64-bit GPRs
- Embedded scalar single-precision floating-point, which use extended 32-bit GPRs
- Embedded scalar double-precision floating-point, which use extended 64-bit GPRs

Note that for 32-bit implementations, the SPE load and store instructions must be used for accessing 64-bit embedded floating-point operands.

### 3.2 SPE Instruction Set

This section describes the data formats and instruction syntax, and provides an overview of computational operations of the SPE instructions.

Chapter 5, "Instruction Set," gives complete descriptions of individual SPE and embedded floating-point instructions.

Opcodes are listed in Appendix B, "SPE and Embedded Floating-Point Opcode Listings."

### 3.2.1 SPE Data Formats

SPE provides integer and fractional data formats, which can be treated as signed or unsigned quantities.

### 3.2.1.1 Integer Format

Unsigned integers consist of 16-, 32-, or 64-bit binary integer values. The largest representable value is $2^{n}-1$, where $n$ represents the number of bits in the value. The smallest representable value is 0 . Computations that produce values larger than $2^{n}-1$ or smaller than 0 set OV or OVH in SPEFSCR.

Signed integers consist of $16-, 32$-, or 64 -bit binary values in two's-complement form. The largest representable value is $2^{n-1}-1$, where $n$ represents the number of bits in the value. The smallest representable value is $-2^{n-1}$. Computations that produce values larger than $2^{n-1}-1$ or smaller than $-2^{n-1}$ set OV or OVH in SPEFSCR.

### 3.2.1.2 Fractional Format

Fractional data is useful for representing data converted from analog devices and is conventionally used for DSP fractional arithmetic.

Unsigned fractions consist of 16-, 32-, or 64-bit binary fractional values that range from 0 to less than 1. Unsigned fractions place the radix point immediately to the left of the msb. The msb of the value represents the value $2^{-1}$, the next msb represents the value $2^{-2}$, and so on. The largest representable value is $1-2^{-n}$ where $n$ represents the number of bits in the value. The smallest representable value is 0 . Computations that produce values larger than $1-2^{-n}$ or smaller than 0 may set OV or OVH in the SPEFSCR. SPE does not define unsigned fractional forms of instructions to manipulate unsigned fractional data because the unsigned integer forms of the instructions produce the same results as unsigned fractional forms.

Guarded unsigned fractions are 64-bit binary fractional values. Guarded unsigned fractions place the decimal point immediately to the left of bit 32 . The largest representable value is $2^{32}-2^{-32}$; the smallest is 0 . Guarded unsigned fractional computations are always modulo and do not set OV or OVH.

Signed fractions consist of 16-, 32-, or 64-bit binary fractional values in two's-complement form that range from -1 to less than 1 . Signed fractions place the decimal point immediately to the right of the msb . The largest representable value is $1-2^{-(n-1)}$ where $n$ represents the number of bits in the value. The smallest representable value is -1 . Computations that produce values larger than $1-2^{-(n-1)}$ or smaller than -1 may set OV or OVH. Multiplication of two signed fractional values causes the result to be shifted left one bit to remove the resultant redundant sign bit in the product. In this case, a 0 bit is concatenated as the lsb of the shifted result.

Guarded signed fractions are 64-bit binary fractional values that place the decimal point immediately to the left of bit 33 . The largest representable value is $2^{32}-2^{-31}$; the smallest is $-2^{32}-1+2^{-31}$. Guarded signed fractional computations are always modulo and do not set OV or OVH.

### 3.2.2 Computational Operations

SPE supports several different computational capabilities. Modulo results produce truncation of the overflow bits in a calculation; therefore, overflow does not occur and no saturation is performed. For instructions for which overflow occurs, saturation provides a maximum or minimum representable value
(for the data type) in the case of overflow. Instructions are provided for a wide range of computational capability. The operation types are as follows:

- Simple vector instructions. These instructions use the corresponding low- and high-word elements of the operands to produce a vector result that is placed in the destination register, the accumulator, or both. Figure 3-1 shows how operations are typically performed in vector operations.


Figure 3-1. Two-Element Vector Operations

- Multiply and accumulate instructions. These instructions perform multiply operations, optionally add the result to the ACC, and place the result into the destination register and optionally into the ACC. These instructions are composed of different multiply forms, data formats, and data accumulate options, as indicated by their mnemonics, as shown in Table 3-1.

Table 3-1. Mnemonic Extensions for Multiply Accumulate Instructions

| Extension | Meaning | Comments |
| :---: | :---: | :---: |
| Multiply Form |  |  |
| he | Half word even | $16 \times 16 \rightarrow 32$ |
| heg | Half word even guarded | $16 \times 16 \rightarrow 32,64-$ bit final accum result |
| ho | Half word odd | $16 \times 16 \rightarrow 32$ |
| hog | Half word odd guarded | $16 \times 16 \rightarrow 32$, 64-bit final accum result |
| w | Word | $32 \times 32 \rightarrow 64$ |
| wh | Word high | $32 \times 32 \rightarrow 32$ (high order 32 bits of product) |
| wl | Word low | $32 \times 32 \rightarrow 32$ (low order 32 bits of product) |
| Data Format |  |  |
| smf | Signed modulo fractional | Modulo, no saturation or overflow |
| smi | Signed modulo integer | Modulo, no saturation or overflow |
| ssf | Signed saturate fractional | Saturation on product and accumulate |
| ssi | Signed saturate integer | Saturation on product and accumulate |
| umi | Unsigned modulo integer | Modulo, no saturation or overflow |
| usi | Unsigned saturate integer | Saturation on product and accumulate |
| Accumulate Option |  |  |
| a | Place in accumulator | Result $\rightarrow$ accumulator |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## SPE and Embedded Floating-Point Instruction Model

Table 3-1. Mnemonic Extensions for Multiply Accumulate Instructions (continued)

| Extension | Meaning | Comments |
| :---: | :---: | :---: |
| aa | Add to accumulator | Accumulator + result $\rightarrow$ accumulator |
| aaw | Add to accumulator | Accumulator $_{0: 31}+$ result $_{0: 31} \rightarrow$ accumulator $_{0: 31}$ <br> Accumulator $_{32: 63}+$ result $_{32: 63} \rightarrow$ accumulator $_{32: 63}$ |
| an | Add negated to accumulator | Accumulator - result $\rightarrow$ accumulator |
| anw | Add negated to accumulator | Accumulator $_{0: 31}$ - result $_{0: 31} \rightarrow$ accumulator $_{0: 31}$ Accumulator $_{32: 63}-$ result $_{32: 63} \rightarrow$ accumulator $_{32: 63}$ |

- Load and store instructions. These instructions provide load and store capabilities for moving data to and from memory. A variety of forms are provided that position data for efficient computation.
- Compare and miscellaneous instructions. These instructions perform miscellaneous functions such as field manipulation, bit reversed incrementing, and vector compares.


### 3.2.2.1 Data Formats and Register Usage

Figure 2-4 shows how GPRs are used with integer, fractional, and floating-point data formats.

### 3.2.2.1.1 Signed Fractions

In signed fractional format, the $n$-bit operand is represented in a 1.[ $n-1]$ format ( 1 sign bit, $n-1$ fraction bits). Signed fractional numbers are in the following range:

$$
-1.0 \leq \mathrm{SF} \leq 1.0-2^{-(n-1)}
$$

The real value of the binary operand $\mathrm{SF}[0: n-1]$ is as follows:

$$
\mathrm{SF}=-1.0 \bullet \mathrm{SF}(0)+\sum_{\mathrm{i}=1}^{n-1} \mathrm{SF}(\mathrm{i}) \bullet 2^{-\mathrm{i}}
$$

The most negative and positive numbers representable in fractional format are as follows:

- The most negative number is represented by $\mathrm{SF}(0)=1$ and $\mathrm{SF}[1: n-1]=0$ (that is, $n=32$; $0 x 8000 \_0000=-1.0$ ).
- The most positive number is represented by $\mathrm{SF}(0)=0$ and $\mathrm{SF}[1: n-1]=$ all 1 s (that is, $n=32$;
$0 \times 7$ FFF_FFFF $=1.0-2^{-(n-1)}$ ).


### 3.2.2.1.2 SPE Integer and Fractional Operations

Figure 3-2 shows data formats for signed integer and fractional multiplication. Note that low word versions of signed saturate and signed modulo fractional instructions are not supported. Attempting to execute an opcode corresponding to these instructions causes boundedly undefined results.


Figure 3-2. Integer and Fractional Operations

### 3.2.2.1.3 SPE Instructions

Table 3-2 shows how SPE vector multiply instruction mnemonics are structured.
Table 3-2. SPE Vector Multiply Instruction Mnemonic Structure

| Prefix | Multiply Element |  | Data Type Element |  | Accumulate Element |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| evm | ho he hog heg wh wl whg wig w | half odd (16x16->32) half even ( $16 \times 16->32$ ) half odd guarded ( $16 \times 16->32$ ) half even guarded (16x16->32) word high (32x32->32) word low (32x32->32) word high guarded (32x32->32) word low guarded (32x32->32) word (32x32->64) | usi umi ssi $\mathbf{s s f}{ }^{1}$ smi $\mathbf{s m f}^{1}$ | unsigned saturate integer unsigned modulo integer signed saturate integer signed saturate fractional signed modulo integer signed modulo fractional | a <br> aa <br> an <br> aaw <br> anw | write to ACC <br> write to ACC \& added ACC <br> write to ACC \& negate ACC <br> write to ACC \& ACC in words <br> write to ACC \& negate ACC in words |

1 Low word versions of signed saturate and signed modulo fractional instructions are not supported. Attempting to execute an opcode corresponding to these instructions causes boundedly undefined results.

Table 3-3 defines mnemonic extensions for these instructions.
Table 3-3. Mnemonic Extensions for Multiply-Accumulate Instructions

| Extension | Meaning |  |  |  |
| :---: | :--- | :--- | :---: | :---: |
| Multiply Form |  |  |  |  |
| he | Half word even | $16 \times 16 \rightarrow 32$ |  |  |
| heg | Half word even guarded | $16 \times 16 \rightarrow 32,64$-bit final accumulator result |  |  |
| ho | Half word odd | $16 \times 16 \rightarrow 32$ |  |  |
| hog | Half word odd guarded | $16 \times 16 \rightarrow 32,64$-bit final accumulator result |  |  |
| w | Word | $32 \times 32 \rightarrow 64$ |  |  |
| wh | Word high | $32 \times 32 \rightarrow 32$, high-order 32 bits of product |  |  |
| wl | Word low | $32 \times 32 \rightarrow 32$, low-order 32 bits of product |  |  |
|  |  |  |  | Data Type |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 3-3. Mnemonic Extensions for Multiply-Accumulate Instructions (continued)

| Extension | Meaning | Comments |  |
| :---: | :--- | :--- | :---: |
| $\mathbf{s m f}$ | Signed modulo fractional | Wrap, no saturate |  |
| $\mathbf{s m i}$ | Signed modulo integer | Wrap, no saturate |  |
| $\mathbf{s s f}$ | Signed saturate fractional | - |  |
| $\mathbf{s s i}$ | Signed saturate integer | - |  |
| $\mathbf{u m i}$ | Unsigned modulo integer | Wrap, no saturate |  |
| $\mathbf{u s i}$ | Unsigned saturate integer | - |  |
| Accumulate Options |  |  |  |
| $\mathbf{a}$ | Update accumulator | Update accumulator (no add) |  |
| aa | Add to accumulator | Add result to accumulator (64-bit sum) |  |
| aaw | Add to accumulator (words) | Add word results to accumulator words (pair of 32-bit sums) |  |
| an | Add negated | Add negated result to accumulator (64-bit sum) |  |
| anw | Add negated to accumulator (words) | Add negated word results to accumulator words (pair of 32-bit sums) |  |

Table 3-4 lists SPE instructions.
Table 3-4. SPE Instructions

| Instruction | Mnemonic | Syntax |
| :--- | :---: | :---: |
| Bit Reversed Increment | brinc | rD,rA,rB |
| Initialize Accumulator | evmra | rD,rA |
| Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate | evmhegsmfan | rD,rA,rB |
| Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate Negative | evmhegsmiaa | rD,rA,rB |
| Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate | evmhegsmian | rD,rA,rB |
| Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate Negative | evmhegumiaa | rD,rA,rB |
| Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate | evmhegumian | rD,rA,rB |
| Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate Negative | evmhogsmfaa | rD,rA,rB |
| Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate | evmhogsmfan | rD,rA,rB |
| Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate Negative | evmhogsmiaa | rD,rA,rB |
| Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer and Accumulate | evmhogsmian | rD,rA,rB |
| Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer and Accumulate Negative | evmhogumian | rD,rA,rB |
| Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate | evabs | rD,rA |
| Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate Negative | evaddiw | rD,rB,UIMM |
| Vector Absolute Value | evaddsmiaaw | rD,rA,rB |
| Vector Add Immediate Word | evaddssiaaw | rD,rA |
| Vector Add Signed, Modulo, Integer to Accumulator Word | evaddumiaaw | rD,rA |
| Vector Add Signed, Saturate, Integer to Accumulator Word | evaddusiaaw | rD,rA |
| Vector Add Unsigned, Modulo, Integer to Accumulator Word | rD,rA,rB |  |
| Vector Add Unsigned, Saturate, Integer to Accumulator Word |  |  |
| Vector Add Word |  |  |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 3-4. SPE Instructions (continued)

| Instruction | Mnemonic | Syntax |
| :---: | :---: | :---: |
| Vector AND | evand | rD,rA,rB |
| Vector AND with Complement | evandc | rD,rA,rB |
| Vector Compare Equal | evcmpeq | crD, rA,rB |
| Vector Compare Greater Than Signed | evcmpgts | crD, rA,rB |
| Vector Compare Greater Than Unsigned | evcmpgtu | crD, rA,rB |
| Vector Compare Less Than Signed | evcmplts | crD,rA,rB |
| Vector Compare Less Than Unsigned | evcmpltu | crD,rA,rB |
| Vector Count Leading Sign Bits Word | evcntlsw | rD,rA |
| Vector Count Leading Zeros Word | evcntlzw | rD,rA |
| Vector Divide Word Signed | evdivws | rD,rA,rB |
| Vector Divide Word Unsigned | evdivwu | rD,rA,rB |
| Vector Equivalent | eveqv | rD,rA,rB |
| Vector Extend Sign Byte | evextsb | rD,rA |
| Vector Extend Sign Half Word | evextsh | rD,rA |
| Vector Load Double into Half Words | evidh | rD,d(rA) |
| Vector Load Double into Half Words Indexed | evldhx | rD,rA,rB |
| Vector Load Double into Two Words | evidw | rD,d(rA) |
| Vector Load Double into Two Words Indexed | evidwx | rD,rA,rB |
| Vector Load Double Word into Double Word | evidd | rD,d(rA) |
| Vector Load Double Word into Double Word Indexed | eviddx | rD,rA,rB |
| Vector Load Half Word into Half Word Odd Signed and Splat | evihhossplat | rD,d(rA) |
| Vector Load Half Word into Half Word Odd Signed and Splat Indexed | evihhossplatx | rD,rA,rB |
| Vector Load Half Word into Half Word Odd Unsigned and Splat | evlhhousplat | rD,d(rA) |
| Vector Load Half Word into Half Word Odd Unsigned and Splat Indexed | evlhhousplatx | rD,rA,rB |
| Vector Load Half Word into Half Words Even and Splat | evlhhesplat | rD,d(rA) |
| Vector Load Half Word into Half Words Even and Splat Indexed | evlhhesplatx | rD,rA,rB |
| Vector Load Word into Half Words and Splat | evlwhsplat | rD,d(rA) |
| Vector Load Word into Half Words and Splat Indexed | evlwhsplatx | rD,rA,rB |
| Vector Load Word into Half Words Odd Signed (with sign extension) | evlwhos | rD,d(rA) |
| Vector Load Word into Half Words Odd Signed Indexed (with sign extension) | evlwhosx | rD,rA,rB |
| Vector Load Word into Two Half Words Even | evlwhe | rD,d(rA) |
| Vector Load Word into Two Half Words Even Indexed | evlwhex | rD,rA,rB |
| Vector Load Word into Two Half Words Odd Unsigned (zero-extended) | evlwhou | rD,d(rA) |
| Vector Load Word into Two Half Words Odd Unsigned Indexed (zero-extended) | evlwhoux | rD,rA,rB |
| Vector Load Word into Word and Splat | evlwwsplat | rD,d(rA) |
| Vector Load Word into Word and Splat Indexed | evlwwsplatx | rD,rA,rB |
| Vector Merge High | evmergehi | rD,rA,rB |
| Vector Merge High/Low | evmergehilo | rD,rA,rB |
| Vector Merge Low | evmergelo | rD,rA,rB |
| Vector Merge Low/High | evmergelohi | rD,rA,rB |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

SPE and Embedded Floating-Point Instruction Model
Table 3-4. SPE Instructions (continued)

| Instruction | Mnemonic | Syntax |
| :---: | :---: | :---: |
| Vector Multiply Half Words, Even, Signed, Modulo, Fractional | evmhesmf | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Fractional and Accumulate into Words | evmhesmfaaw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Fractional and Accumulate Negative into Words | evmhesmfanw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Fractional, Accumulate | evmhesmfa | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Integer | evmhesmi | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate into Words | evmhesmiaaw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate Negative into Words | evmhesmianw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Modulo, Integer, Accumulate | evmhesmia | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Fractional | evmhessf | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate into Words | evmhessfaaw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate Negative into Words | evmhessfanw | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Fractional, Accumulate | evmhessfa | rD,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate into Words | evmhessiaaw | D,rA,rB |
| Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate Negative into Words | evmhessianw | rD,rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Modulo, Integer | evmheumi | rD,rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate into Words | evmheumiaaw | rD, rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate Negative into Words | evmheumianw | rD, rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Modulo, Integer, Accumulate | evmheumia | rD,rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate into Words | evmheusiaaw | rD,rA,rB |
| Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate Negative into Words | evmheusianw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Fractional | evmhosmf | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate into Words | evmhosmfaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate Negative into Words | evmhosmfanw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Fractional, Accumulate | evmhosmfa | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Integer | evmhosmi | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate into Words | evmhosmiaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate Negative into Words | evmhosmianw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Modulo, Integer, Accumulate | evmhosmia | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Fractional | evmhossf | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate into Words | evmhossfaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate Negative into Words | evmhossfanw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Fractional, Accumulate | evmhossfa | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate into Words | evmhossiaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate Negative into Words | evmhossianw | rD,rA,rB |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 3-4. SPE Instructions (continued)

| Instruction | Mnemonic | Syntax |
| :---: | :---: | :---: |
| Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer | evmhoumi | rD,rA,rB |
| Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate into Words | evmhoumiaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate Negative into Words | evmhoumianw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer, Accumulate | evmhoumia | rD,rA,rB |
| Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate into Words | evmhousiaaw | rD,rA,rB |
| Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate Negative into Words | evmhousianw | rD,rA,rB |
| Vector Multiply Word High Signed, Modulo, Fractional | evmwhsmf | rD,rA,rB |
| Vector Multiply Word High Signed, Modulo, Fractional and Accumulate | evmwhsmfa | rD,rA,rB |
| Vector Multiply Word High Signed, Modulo, Integer | evmwhsmi | rD,rA,rB |
| Vector Multiply Word High Signed, Modulo, Integer and Accumulate | evmwhsmia | rD,rA,rB |
| Vector Multiply Word High Signed, Saturate, Fractional | evmwhssf | rD,rA,rB |
| Vector Multiply Word High Signed, Saturate, Fractional and Accumulate | evmwhssfa | rD,rA,rB |
| Vector Multiply Word High Unsigned, Modulo, Integer | evmwhumi | rD,rA,rB |
| Vector Multiply Word High Unsigned, Modulo, Integer and Accumulate | evmwhumia | rD,rA,rB |
| Vector Multiply Word Low Signed, Modulo, Integer and Accumulate in Words | evmwlsmiaaw | rD,rA,rB |
| Vector Multiply Word Low Signed, Modulo, Integer and Accumulate Negative in Words | evmwlsmianw | D,rA,rB |
| Vector Multiply Word Low Signed, Saturate, Integer and Accumulate in Words | evmwIssiaaw | D,rA,rB |
| Vector Multiply Word Low Signed, Saturate, Integer and Accumulate Negative in Words | evmwlssianw | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Modulo, Integer | evmwlumi | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate | evmwlumia | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate in Words | evmwlumiaaw | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate Negative in Words | evmwlumianw | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate in Words | evmwlusiaaw | rD,rA,rB |
| Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate Negative in Words | evmwlusianw | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Fractional | evmwsmf | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Fractional and Accumulate | evmwsmfa | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Fractional and Accumulate | evmwsmfaa | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Fractional and Accumulate Negative | evmwsmfan | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Integer | evmwsmi | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Integer and Accumulate | evmwsmia | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Integer and Accumulate | evmwsmiaa | rD,rA,rB |
| Vector Multiply Word Signed, Modulo, Integer and Accumulate Negative | evmwsmian | rD,rA,rB |
| Vector Multiply Word Signed, Saturate, Fractional | evmwssf | rD,rA,rB |
| Vector Multiply Word Signed, Saturate, Fractional and Accumulate | evmwssfa | rD,rA,rB |
| Vector Multiply Word Signed, Saturate, Fractional and Accumulate | evmwssfaa | rD, rA,rB |
| Vector Multiply Word Signed, Saturate, Fractional and Accumulate Negative | evmwssfan | rD,rA,rB |
| Vector Multiply Word Unsigned, Modulo, Integer | evmwumi | rD,rA,rB |
| Vector Multiply Word Unsigned, Modulo, Integer and Accumulate | evmwumia | rD,rA,rB |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## SPE and Embedded Floating-Point Instruction Model

Table 3-4. SPE Instructions (continued)

| Instruction | Mnemonic | Syntax |
| :---: | :---: | :---: |
| Vector Multiply Word Unsigned, Modulo, Integer and Accumulate | evmwumiaa | rD,rA,rB |
| Vector Multiply Word Unsigned, Modulo, Integer and Accumulate Negative | evmwumian | rD,rA,rB |
| Vector NAND | evnand | rD,rA,rB |
| Vector Negate | evneg | rD,rA |
| Vector NOR ${ }^{1}$ | evnor | rD,rA,rB |
| Vector OR ${ }^{2}$ | evor | rD,rA,rB |
| Vector OR with Complement | evorc | rD,rA,rB |
| Vector Rotate Left Word | evrlw | rD,rA,rB |
| Vector Rotate Left Word Immediate | evrlwi | rD,rA,UIMM |
| Vector Round Word | evrndw | rD,rA |
| Vector Select | evsel | rD,rA,rB,crS |
| Vector Shift Left Word | evslw | rD,rA,rB |
| Vector Shift Left Word Immediate | evslwi | rD,rA,UIMM |
| Vector Shift Right Word Immediate Signed | evsrwis | rD,rA,UIMM |
| Vector Shift Right Word Immediate Unsigned | evsrwiu | rD,rA,UIMM |
| Vector Shift Right Word Signed | evsrws | rD,rA,rB |
| Vector Shift Right Word Unsigned | evsrwu | rD,rA,rB |
| Vector Splat Fractional Immediate | evsplatfi | rD,SIMM |
| Vector Splat Immediate | evsplati | rD,SIMM |
| Vector Store Double of Double | evstdd | rS,d(rA) |
| Vector Store Double of Double Indexed | evstddx | rS,rA,rB |
| Vector Store Double of Four Half Words | evstdh | rS, d(rA) |
| Vector Store Double of Four Half Words Indexed | evstdhx | rS,rA,rB |
| Vector Store Double of Two Words | evstdw | rS,d(rA) |
| Vector Store Double of Two Words Indexed | evstdwx | rS,rA,rB |
| Vector Store Word of Two Half Words from Even | evstwhe | rS, $\mathrm{d}(\mathrm{rA})$ |
| Vector Store Word of Two Half Words from Even Indexed | evstwhex | rS,rA,rB |
| Vector Store Word of Two Half Words from Odd | evstwho | rS, d(rA) |
| Vector Store Word of Two Half Words from Odd Indexed | evstwhox | rS,rA,rB |
| Vector Store Word of Word from Even | evstwwex | rS,d(rA) |
| Vector Store Word of Word from Even Indexed | evstwwex | rS,rA,rB |
| Vector Store Word of Word from Odd | evstwwo | rS,d(rA) |
| Vector Store Word of Word from Odd Indexed | evstwwox | rS,rA,rB |
| Vector Subtract from Word ${ }^{3}$ | evsubfw | rD,rA,rB |
| Vector Subtract Immediate from Word ${ }^{4}$ | evsubifw | rD,UIMM,rB |
| Vector Subtract Signed, Modulo, Integer to Accumulator Word | evsubfsmiaaw | rD,rA |
| Vector Subtract Signed, Saturate, Integer to Accumulator Word | evsubfssiaaw | rD, rA |
| Vector Subtract Unsigned, Modulo, Integer to Accumulator Word | evsubfumiaaw | rD,rA |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 3-4. SPE Instructions (continued)

| Instruction | Mnemonic | Syntax |
| :--- | :---: | :---: |
| Vector Subtract Unsigned, Saturate, Integer to Accumulator Word | evsubfusiaaw | rD,rA |
| Vector XOR | evxor | rD,rA,rB |

1 evnot rD,rA is equivalent to evnor rD,rA,rA
2 evmr rD,rA is equivalent to evor rD,rA,rA
3 evsubw rD,rB,rA is equivalent to evsubfw rD,rA,rB
4 evsubiw rD,rB, UIMM is equivalent to evsubifw $r D$, UIMM, $r B$

### 3.2.3 SPE Simplified Mnemonics

Table 3-5 lists simplified mnemonics for SPE instructions.
Table 3-5. SPE Simplified Mnemonics

| Simplified Mnemonic | Equivalent |
| :---: | :--- |
| evmr rD,rA | evor rD,rA,rA |
| evnot rD,rA | evnor rD,rA,rA |
| evsubiw rD,rB,UIMM | evsubifw rD,UIMM,rB |
| evsubw rD,rB,rA | evsubfw rD,rA,rB |

### 3.3 Embedded Floating-Point Instruction Set

The embedded floating-point categories require the implementation of the signal processing engine (SPE) category and consist of three distinct categories:

- Embedded vector single-precision floating-point
- Embedded scalar single-precision floating-point
- Embedded scalar double-precision floating-point

Although each of these may be implemented independently, they are defined in a single chapter because they may be implemented together.

Load and store instructions for transferring operands to and from memory are described in Section 3.3.3, "Load/Store Instructions."

References to embedded floating-point categories, embedded floating-point instructions, or embedded floating-point operations apply to all three categories.

Scalar single-precision floating-point operations use 32-bit GPRs as source and destination operands; however, double precision and vector instructions require 64-bit GPRs as described in Section 2.2.1, "General-Purpose Registers (GPRs)."

Opcodes are listed in Appendix B, "SPE and Embedded Floating-Point Opcode Listings."

### 3.3.1 Embedded Floating-Point Operations

This section describes embedded floating-point operational modes, data formats, underflow and overflow handling, compliance with IEEE 754, and conversion models.

### 3.3.1.1 Operational Modes

All embedded floating-point operations are governed by the setting of the mode bit in SPEFSCR. The mode bit defines how floating-point results are computed and how floating-point exceptions are handled. Mode 0 defines a real-time, default-results-oriented mode that saturates results. Other modes are currently not defined.

### 3.3.1.2 Floating-Point Data Formats

Single-precision floating-point data elements are 32 bits wide with 1 sign bit (s), 8 bits of biased exponent (e) and 23 bits of fraction (f). Double-precision floating-point data elements are 64 bits wide with 1 sign bit (s), 11 bits of biased exponent (e) and 52 bits of fraction (f).

In the IEEE-754 specification, floating-point values are represented in a format consisting of three explicit fields (sign field, biased exponent field, and fraction field) and an implicit hidden bit. Figure 3-3 shows floating-point data formats.

s-sign bit; $0=$ positive; $1=$ negative
exp-biased exponent field
fraction-fractional portion of number
Figure 3-3. Floating-Point Data Format
For single-precision normalized numbers, the biased exponent value e lies in the range of 1 to 254 corresponding to an actual exponent value E in the range -126 to +127 . For double-precision normalized numbers, the biased exponent value e lies in the range of 1 to 2046 corresponding to an actual exponent value E in the range -1022 to +1023 . With the hidden bit implied to be ' 1 ' (for normalized numbers), the value of the number is interpreted as follows:

$$
(-1)^{s} \times 2^{E} \times(1 . \text { fraction })
$$

where E is the unbiased exponent and 1.fraction is the mantissa (or significand) consisting of a leading ' 1 ' (the hidden bit) and a fractional part (fraction field). For the single-precision format, the maximum positive normalized number (pmax) is represented by the encoding 0x7F7F_FFFF which is approximately $3.4 \mathrm{E}+38,\left(2^{128}\right)$, and the minimum positive normalized value ( pmin ) is represented by the encoding $0 x 0080 \_0000$ which is approximately $1.2 \mathrm{E}-38\left(2^{-126}\right)$. For the double-precision format, the maximum positive normalized number (pmax) is represented by the encoding 0x7FEF_FFFF_FFFF_FFFF which is approximately $1.8 \mathrm{E}+307\left(2^{1024}\right)$, and the minimum positive normalized value ( pmin ) is represented by the encoding 0x0010_0000_0000_0000 which is approximately $2.2 \mathrm{E}-308\left(2^{-1022}\right)$.

Two specific values of the biased exponent are reserved ( 0 and 255 for single-precision; 0 and 2047 for double-precision) for encoding special values of $+0,-0,+$ infinity, -infinity, and NaNs.

Zeros of both positive and negative sign are represented by a biased exponent value e of 0 and a fraction $f$ which is 0 .

Infinities of both positive and negative sign are represented by a maximum exponent field value ( 255 for single-precision, 2047 for double-precision) and a fraction which is 0 .

Denormalized numbers of both positive and negative sign are represented by a biased exponent value e of 0 and a fraction f , which is nonzero. For these numbers, the hidden bit is defined by IEEE 754 to be 0 . This number type is not directly supported in hardware. Instead, either a software interrupt handler is invoked, or a default value is defined.

NaNs (Not-a-Numbers) are represented by a maximum exponent field value ( 255 for single-precision, 2047 for double-precision) and a fraction, f , which is nonzero.

### 3.3.1.3 Overflow and Underflow

Defining pmax to be the most positive normalized value (farthest from zero), pmin the smallest positive normalized value (closest to zero), nmax the most negative normalized value (farthest from zero) and nmin the smallest normalized negative value (closest to zero), an overflow is said to have occurred if the numerically correct result of an instruction is such that $r$ > pmax or $r$ < nmax. Additionally, an implementation may also signal overflow by comparing the exponents of the operands. In this case, the hardware examines both exponents ignoring the fractional values. If it is determined that the operation to be performed may overflow (ignoring the fractional values), an overflow may be said to occur. For addition and subtraction this can occur if the larger exponent of both operands is 254 . For multiplication this can occur if the sum of the exponents of the operands less the bias is 254 . Thus:

```
single-precision addition:
    if A Aexp >= 254 | B Bexp >= 254 then overflow
double-precision addition:
    if A Axp >= 2046 | Bexp >= 2046 then overflow
single-precision multiplication:
    if A Axp + B exp - 127 >= 254 then overflow
double-precision multiplication:
    if A Axp + Bexp - 1023 >= 2046 then overflow
```

An underflow is said to have occurred if the numerically correct result of an instruction is such that $0<\mathrm{r}<$ pmin or $\operatorname{m\operatorname {min}}<\mathrm{r}<0$. In this case, r may be denormalized, or may be smaller than the smallest denormalized number. As with overflow detection, an implementation may also signal underflow by comparing the exponents of the operands. In this case, the hardware examines both exponents regardless of the fractional values. If it is determined that the operation to be performed may underflow (ignoring the fractional values), an underflow may be said to occur. For division, this can occur if the difference of the exponent of the A operand less the exponent of the B operand less the bias is 1 . Thus:

```
single-precision division:
    if A}\mp@subsup{A}{\operatorname{exp}}{}-\mp@subsup{B}{\operatorname{exp}}{}-127<=1 then underflow
double-precision multiplication:
    if A Axp - Bexp - 1023 <= 1 then underflow
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Embedded floating-point operations do not produce + Inf, $-\operatorname{Inf}$, NaN, or a denormalized number. If the result of an instruction overflows and floating-point overflow exceptions are disabled
(SPEFSCR[FOVFE] is cleared), pmax or nmax is generated as the result of that instruction depending on the sign of the result. If the result of an instruction underflows and floating-point underflow exceptions are disabled (SPEFSCR[FUNFE] is cleared), +0 or -0 is generated as the result of that instruction based upon the sign of the result.

### 3.3.1.4 IEEE Std 754 ${ }^{\text {TM }}$ Compliance

The embedded floating-point categories require a floating-point system as defined in IEEE 754 but may rely on software support in order to conform fully with the standard. Thus, whenever an input operand of the embedded floating-point instruction has data values that are +infinity, -infinity, alized, NaN , or when the result of an operation produces an overflow or an underflow, an embedded floating-point data interrupt may be taken and the interrupt handler is responsible for delivering IEEE 754-compliant behavior if desired.

When embedded floating-point invalid operation/input error exceptions are disabled (SPEFSCR[FINVE] $=0$ ), default results are provided by the hardware when an infinity, denormalized, or NaN input is received, or for the operation $0 / 0$. When embedded floating-point underflow exceptions are disabled
(SPEFSCR[FUNFE] $=0$ ) and the result of a floating-point operation underflows, a signed zero result is produced. The embedded floating-point round (inexact) exception is also signaled for this condition. When embedded floating-point overflow exceptions are disabled (SPEFSCR[FOVFE] $=0$ ) and the result of a floating-point operation overflows, a pmax or nmax result is produced. The embedded floating-point round (inexact) exception is also signaled for this condition. An exception enable flag (SPEFSCR[FINXE]) is also provided for generating an embedded floating-point round interrupt when an inexact result is produced, to allow a software handler to conform to IEEE 754. An embedded floating-point divide by zero exception enable flag (SPEFSCR[FDBZE]) is provided for generating an embedded floating-point data interrupt when a divide by zero operation is attempted to allow a software handler to conform to IEEE 754. All of these exceptions may be disabled, and the hardware will then deliver an appropriate default result.

The sign of the result of an addition operation is the sign of the source operand having the larger absolute value. If both operands have the same sign, the sign of the result is the same as the sign of the operands. This includes subtraction which is addition with the negation of the sign of the second operand. The sign of the result of an addition operation with operands of differing signs for which the result is zero is positive except when rounding to negative infinity. Thus $-0+-0=-0$, and all other cases which result in a zero value give +0 unless the rounding mode is rounded to negative infinity.

## NOTE (Programming)

When exceptions are disabled and default results computed, operations having input values that are denormalized may provide different results on different implementations. An implementation may choose to use the denormalized value or a zero value for any computation. Thus a computational operation involving a denormalized value and a normal value may return different results depending on the implementation.

### 3.3.1.5 Sticky Bit Handling for Exception Conditions

The SPEFSCR defines sticky bits for retaining information about exception conditions that are detected. These sticky bits (FINXS, FINVS, FDBZS, FUNFS, and FOVFS) can be used to help provide IEEE-754 compliance. The sticky bits represent the combined OR of all previous status bits produced from any embedded floating-point operation before the last time software zeroed the sticky bit. Only software can zero a sticky bit; hardware can only set sticky bits.

The SPEFSCR is described in Section 2.2.3, "Signal Processing Embedded Floating-Point Status and Control Register (SPEFSCR)." Interrupts are described in Chapter 4, "SPE/Embedded Floating-Point Interrupt Model."

### 3.3.1.6 Implementation Options Summary

There are several options that may be chosen for a given implementation. This section summarizes implementation-dependent functionality and should be used with the processor core documentation to determine behavior of individual implementations.

- Floating-point instruction sets can be implemented independently of one another.
- Overflow and underflow conditions may be signaled by evaluating the exponent. If the evaluaton indicates an overflow or underflow could occur, the implementation may choose to signal an overflow or underflow. It is recommended that future implementations not use this estimation and that they signal overflow or underflow when they actually occur.
- If an operand for a calculation or conversion is denormalized, the implementation may choose to use a same-signed zero value in place of the denormalized operand.
- The rounding modes of +infinity and -infinity are not required to be handled by an implementation. If an implementation does not support $\pm$ infinity rounding modes and the rounding mode is set to be +infinity or -infinity, an embedded floating-point round interrupt occurs after every floating-point instruction for which rounding may occur, regardless of the value of FINXE, unless an embedded floating-point data interrupt also occurs and is taken.
- For absolute value, negate, and negative absolute value operations, an implementation may choose either to simply perform the sign bit operation, ignoring exceptions, or to compute the operation and handle exceptions and saturation where appropriate.
- SPEFSCR[FGH,FXH] are undefined on completion of a scalar floating-point operation. An implementation may choose to zero them or leave them unchanged.
- An implementation may choose to only implement sticky bit setting by hardware for FDBZS and FINXS, allowing software to manage the other sticky bits. It is recommended that all future implementations implement all sticky bit setting in hardware.
- For 64-bit implementations, the upper 32 bits of the destination register are undefined when the result of a scalar floating-point operation is a 32-bit result. It is recommended that future 64-bit implementations produce 64 -bit results for the results of 64 -bit convert-to-integer values.


### 3.3.1.7 Saturation, Shift, and Bit Reverse Models

For saturation, left shifts, and bit reversal, the pseudo-RTL is provided here to more accurately describe those functions referenced in the instruction pseudo-RTL.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 3.3.1.7.1 Saturation

```
SATURATE(ov, carry, sat_ovn, sat_ov, val)
if ov then
    if carry then
        return sat_ovn
    else
        return sat_ov
else
    return val
```


### 3.3.1.7.2 Shift Left

```
SL(value, cnt)
if cnt > 31 then
    return 0
else
    return (value << cnt)
```


### 3.3.1.7.3 Bit Reverse

```
BITREVERSE (value)
result \leftarrow0
mask \leftarrow1
shift \leftarrow < 31
cnt }\leftarrow3
while cnt > 0 then do
    t \leftarrowvalue & mask
    if shift >= 0 then
        result \leftarrow(t << shift) | result
    else
        result \leftarrow(t >> -shift) | result
    cnt \leftarrowcnt - 1
    shift \leftarrowshift - 2
    mask \leftarrowmask << 1
return result
```


### 3.3.2 Embedded Vector and Scalar Floating-Point Instructions

The embedded floating-point operations are IEEE 754-compliant with software exception handlers and offer a simpler exception model than the Power ISA floating-point instructions that use the floating-point registers (FPRs). Instead of FPRs, these instructions use GPRs to offer improved performance for converting among floating-point, integer, and fractional values. Sharing GPRs allows vector floating-point instructions to use SPE load and store instructions.

## NOTE

Note that the vector and scalar versions of the instructions have the same syntax.

Table 3-6 lists the vector and scalar floating-point instructions.
Table 3-6. Vector and Scalar Floating-Point Instructions

| Instruction | Single-Precision |  | DoublePrecision Scalar | Syntax |
| :---: | :---: | :---: | :---: | :---: |
|  | Scalar | Vector |  |  |
| Convert Floating-Point Double- from Single-Precision | - | - | efdcfs | rD, rB |
| Convert Floating-Point from Signed Fraction | efscfsf | evfscfsf | efdcfsf | rD, rB |
| Convert Floating-Point from Signed Integer | efscfsi | evfscfsi | efdcfsi | rD,rB |
| Convert Floating-Point from Unsigned Fraction | efscfuf | evfscfuf | efdcfuf | rD, rB |
| Convert Floating-Point from Unsigned Integer | efscfui | evfscfui | efdcfui | rD,rB |
| Convert Floating-Point Single- from Double-Precision | - | - | efscfd | rD,rB |
| Convert Floating-Point to Signed Fraction | efsctsf | evfsctsf | efdctsf | rD, rB |
| Convert Floating-Point to Signed Integer | efsctsi | evfsctsi | efdctsi | rD,rB |
| Convert Floating-Point to Signed Integer with Round toward Zero | efsctsiz | evfsctsiz | efdctsiz | rD,rB |
| Convert Floating-Point to Unsigned Fraction | efsctuf | evfsctuf | efdctuf | rD,rB |
| Convert Floating-Point to Unsigned Integer | efsctui | evfsctui | efdctui | rD, rB |
| Convert Floating-Point to Unsigned Integer with Round toward Zero | efsctuiz | evfsctuiz | efdctuiz | rD,rB |
| Floating-Point Absolute Value | efsabs ${ }^{1}$ | evfsabs | efdabs | rD, rA |
| Floating-Point Add | efsadd | evfsadd | efdadd | rD,rA,rB |
| Floating-Point Compare Equal | efscmpeq | evfscmpeq | efdcmpeq | crD,rA,rB |
| Floating-Point Compare Greater Than | efscmpgt | evfscmpgt | efdcmpgt | crD,rA,rB |
| Floating-Point Compare Less Than | efscmplt | evfscmplt | efdcmplt | crD,rA,rB |
| Floating-Point Divide | efsdiv | evfsdiv | efddiv | rD,rA,rB |
| Floating-Point Multiply | efsmul | evfsmul | efdmul | rD,rA,rB |
| Floating-Point Negate | efsneg ${ }^{1}$ | evfsneg | efdneg | rD, rA |
| Floating-Point Negative Absolute Value | efsnabs ${ }^{1}$ | evfsnabs | efdnabs | rD, rA |
| Floating-Point Subtract | efssub | evfssub | efdsub | rD,rA,rB |
| Floating-Point Test Equal | efststeq | evfststeq | efdtsteq | crD,rA,rB |
| Floating-Point Test Greater Than | efststgt | evfststgt | efdtstgt | crD,rA,rB |
| Floating-Point Test Less Than | efststlt | evfststlt | efdtstlt | crD,rA,rB |
| SPE Double Word Load/Store Instructions |  |  |  |  |
| Vector Load Double Word into Double Word | - | evidd | evidd | rD, d(rA) |
| Vector Load Double Word into Double Word Indexed | - | eviddx | eviddx | rD,rA,rB |
| Vector Merge High | - | evmergehi | evmergehi | rD,rA,rB |
| Vector Merge Low | - | evmergelo | evmergelo | rD,rA,rB |
| Vector Store Double of Double | - | evstdd | evstdd | rS,d(rA) |
| Vector Store Double of Double Indexed | - | evstddx | evstddx | rS,rA,rB |

Note: On some cores, floating-point operations that produce a result of zero may generate an incorrect sign.
1 Exception detection for these instructions is implementation dependent. On some devices, infinities, NaNs , and denorms are always be treated as Norms. No exceptions are taken if SPEFSCR[FINVE] = 1.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 3.3.3 Load/Store Instructions

Embedded floating-point instructions use GPRs to hold and operate on floating-point values. Standard load and store instructions are used to move the data to and from memory. If vector single-precision or scalar double-precision embedded floating-point instructions are implemented on a 32-bit implementation, the GPRs are 64 bits wide. Because a 32-bit implementation contains no load or store instructions that operate on 64-bit data, the following SPE load/store instructions are used:

- evldd-Vector Load Doubleword into Doubleword
- evlddx—Vector Load Doubleword into Doubleword Indexed
- evstdd—Vector Store Doubleword of Doubleword
- evstddx-Vector Store Doubleword of Doubleword
- evmergehi-Vector Merge High
- evmergelo-Vector Merge Low


### 3.3.3.1 Floating-Point Conversion Models

Pseudo-RTL models for converting floating-point to and from non-floating-point is provided in Section 5.3.2, "Embedded Floating-Point Conversion Models," as a group of functions called from the individual instruction pseudo-RTL descriptions, which are included in the instruction descriptions in Chapter 5, "Instruction Set."

## Chapter 4 <br> SPE/Embedded Floating-Point Interrupt Model

This chapter describes the SPE interrupt model, including the SPE embedded floating-point interrupts

### 4.1 Overview

The SPE defines additional exceptions that can generate an alignment interrupt and three additional interrupts to allow software handling of exceptions that may occur during execution of SPE.embedded floating-point instructions. These are shown in Table 4-1 and described in detail in the following sections.

Table 4-1. SPE/SPE Embedded Floating-Point Interrupt and Exception Types

| IVOR | Interrupt | Exception | Synchronous/ Precise | ESR | MSR <br> Mask | DBCR0/TCR Mask | Category | Page |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| IVOR5 | Alignment | Alignment | Synchronous/ Precise | [ST],[FP,AP,SPV] [EPID],[VLEMI] | - | - | SPE/ <br> Embedded FP | 4.2.2/4-2 |
| IVOR32 | SPE/embedded floating-point ${ }^{1}$ | SPE unavailable | Synchronous/ Precise | SPV, [VLEMI] | - | - | SPE | 4.2.3/4-2 |
| IVOR33 | Embedded floating-point data | Embedded floating-point data | Synchronous/ Precise | SPV, [VLEMI] | - | - | Embedded FP | 4.2.4/4-3 |
| IVOR34 | Embedded floating-point round | Embedded floating-point round | Synchronous/ Precise | SPV, [VLEMI] | - | - | Embedded FP | 4.2.2/4-2 |

1 Other implementations use IVOR32 for vector (AltiVec) unavailable interrupts.

### 4.2 SPE Interrupts

This section describes the interrupts that can be generated when an SPE/embedded floating-point exception is encountered.

### 4.2.1 Interrupt-Related Registers

Figure 4-1 shows the register resources that are defined by the base category and by the SPE interrupt model. Base category resources are described in the EREF.


Figure 4-1. SPE Interrupt-Related Registers

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0


Figure 4-1. SPE Interrupt-Related Registers

### 4.2.2 Alignment Interrupt

An SPE vector alignment exception occurs if the EA of any of the following instructions in not aligned to a 64-bit boundary: evldd, evlddx, evldw, evldwx, evldh, evldhx, evstdd, evstddx, evstdw, evstdwx, evstdh, or evstdhx. When an SPE vector alignment exception occurs, an alignment interrupt is taken and the processor suppresses execution of the instruction causing the exception. SRR0, SRR1, MSR, ESR, and DEAR are modified as follows:

- SRR0 is set to the EA of the instruction causing the interrupt.
- SRR1 is set to the contents of the MSR at the time of the interrupt.
- MSR bits CE, ME, and DE are unchanged. All other bits are cleared.
- ESR[SPV] is set. ESR[ST] is set if the instruction causing the interrupt is a store. All other ESR bits are cleared.
- DEAR is updated with the EA of the access that caused the exception. This is generally the EA of the instruction, except for some instructions that are misaligned or that reference multiple storage element.

Instruction execution resumes at address IVPR[0-47]||IVOR5[48-59]||0b0000.

### 4.2.3 SPE/Embedded Floating-Point Unavailable Interrupt

An SPE/embedded floating-point unavailable exception occurs on an attempt to execute any of the following instructions and MSR[SPV] is not set:

- SPE instruction (except brinc)
- An embedded scalar double-precision instruction
- A vector single-precision floating-point instructions

It is not used by embedded scalar single-precision floating-point instructions.
If this exception occurs, an SPE/embedded floating-point unavailable interrupt is taken and the processor suppresses execution of the instruction causing the exception. Registers are modified as follows:

The SRR0, SRR1, MSR, and ESR registers are modified as follows:

- SRR0 is set to the EA of the instruction causing the interrupt.
- SRR1 is set to the contents of the MSR at the time of the interrupt.
- MSR bits CE, ME, and DE are unchanged. All other bits are cleared.
- ESR bits SPV (and VLEMI if VLE is implemented and the instruction causing the interrupt resides in VLE storage) are set. All other ESR bits are cleared.

Instruction execution resumes at address IVPR[0-47]||IVOR32[48-59]||0b0000.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## NOTE (Software)

Software should use this interrupt to determine if the application is using the upper 32 bits of the GPRs and thus is required to save and restore them on a context switch.

### 4.2.4 SPE Embedded Floating-Point Interrupts

The following sections describe SPE embedded floating-point interrupts:

- Section 4.2.4.1, "Embedded Floating-Point Data Interrupt"
- Section 4.2.4.2, "Embedded Floating-Point Round Interrupt"


### 4.2.4.1 Embedded Floating-Point Data Interrupt

The embedded floating-point data interrupt vector is used for enabled floating-point invalid operation/input error, underflow, overflow, and divide-by-zero exceptions (collectively called floating-point data exceptions). When one of these enabled exceptions occurs, the processor suppresses execution of the instruction causing the exception. The SRR0, SRR1, MSR, ESR, and SPEFSCR are modified as follows:

- SRR0 is set to the EA of the instruction causing the interrupt.
- SRR1 is set to the contents of the MSR at the time of the interrupt.
- MSR bits CE, ME, and DE are unchanged. All other bits are cleared.
- ESR[SPV] is set. All other ESR bits are cleared.
- One or more SPEFSCR status bits are set to indicate the type of exception. The affected bits are FINVH, FINV, FDBZH, FDBZ, FOVFH, FOVF, FUNFH, and FUNF. SPEFSCR[FG,FGH, FX, FXH] are cleared.

Instruction execution resumes at address IVPR[0-47]||IVOR33[48-59]||0b0000.

### 4.2.4.2 Embedded Floating-Point Round Interrupt

The embedded floating-point round interrupt occurs if no other floating-point data interrupt is taken and one of the following conditions is met:

- SPEFSCR[FINXE] is set and the unrounded result of an operation is not exact
- SPEFSCR[FINXE] is set, an overflow occurs, and overflow exceptions are disabled (FOVF or FOVFH set with FOVFE cleared)
- An underflow occurs and underflow exceptions are disabled (FUNF set with FUNFE cleared)

The embedded floating-point round interrupt does not occur if an enabled embedded floating-point data interrupt occurs.

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## NOTE (Programming)

If an implementation does not support $\pm$ infinity rounding modes and the rounding mode is set to be +infinity or -infinity, an embedded floating-point round interrupt occurs after every embedded floating-point instruction for which rounding might occur regardless of the FINXE value, if no higher priority exception exists.

When an embedded floating-point round interrupt occurs, the unrounded (truncated) result of an inexact high or low element is placed in the target register. If only a single element is inexact, the other exact element is updated with the correctly rounded result, and the FG and FX bits corresponding to the other exact element are be 0 .

FG (FGH) and FX (FXH) are provided so an interrupt handler can round the result as it desires. $\mathrm{FG}(\mathrm{FGH})$ is the value of the bit immediately to the right of the lsb of the destination format mantissa from the infinitely precise intermediate calculation before rounding. FX (FXH) is the value of the OR of all bits to the right of the FG (FGH) of the destination format mantissa from the infinitely precise intermediate calculation before rounding.

The SRR0, SRR1, MSR, ESR, and SPEFSCR are modified as follows:

- SRR0 is set to the EA of the instruction following the instruction causing the interrupt.
- SRR1 is set to the contents of the MSR at the time of the interrupt.
- MSR bits CE, ME, and DE are unchanged. All other bits are cleared.
- ESR[SPV] is set. All other ESR bits are cleared.
- SPEFSCR[FGH,FG,FXH,FX] are set appropriately. SPEFSCR[FINXS] is set.

Instruction execution resumes at address IVPR[0-47]||IVOR34[48-59]||0b0000.

### 4.3 Interrupt Priorities

The priority order among the SPE and embedded floating-point interrupts is as follows:

1. SPE/embedded floating-point unavailable interrupt
2. SPE vector alignment interrupt
3. Embedded floating-point data interrupt
4. Embedded floating-point round interrupt

The EREF describes how these interrupts are prioritized among the other Power ISA interrupts. Only one of the above types of synchronous interrupts may have an existing exception generating it at any given time. This is guaranteed by the exception priority mechanism and the requirements of the sequential execution model.

### 4.4 Exception Conditions

The following sections describe the exception conditions that can generate the interrupts described in Section 4.2, "SPE Interrupts." Enable and status bits associated with these programming exceptions can

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0
be found in the SPEFSCR, described in Section 2.2.3, "Signal Processing Embedded Floating-Point Status and Control Register (SPEFSCR)."

### 4.4.1 Floating-Point Exception Conditions

This section describes the conditions that generate exceptions that, depending on how the processor is configured, may generate an interrupt.

### 4.4.1.1 Denormalized Values on Input

Any denormalized value used as an operand may be truncated by the implementation to a properly signed zero value.

### 4.4.1.2 Embedded Floating-Point Overflow and Underflow

Defining pmax to be the most positive normalized value (farthest from zero), pmin the smallest positive normalized value (closest to zero), nmax the most negative normalized value (farthest from zero) and nmin the smallest normalized negative value (closest to zero), an overflow is said to have occurred if the numerically correct result ( r ) of an instruction is such that $\mathrm{r}>$ pmax or $\mathrm{r}<n \mathrm{nmax}$. An underflow is said to have occurred if the numerically correct result of an instruction is such that $0<r<p m i n$ or $n m i n<r<0$. In this case, $r$ may be denormalized, or may be smaller than the smallest denormalized number.

The embedded floating-point categories do not produce +infinity, -infinity, NaN, or denormalized numbers. If the result of an instruction overflows and embedded floating-point overflow exceptions are disabled (SPEFSCR[FOVFE]=0), pmax or nmax is generated as the result of that instruction depending upon the sign of the result. If the result of an instruction underflows and embedded floating-point underflow exceptions are disabled (SPEFSCR[FUNFE]=0), +0 or -0 is generated as the result of that instruction based upon the sign of the result.

If an overflow occurs, SPEFSCR[FOVF FOVFH] are set appropriately, or if an underflow occurs, SPEFSCR[FUNF FUNFH] are set appropriately. If either embedded floating-point underflow or embedded floating-point overflow exceptions are enabled and a corresponding status bit is 1 , an embedded floating-point data interrupt is taken and the destination register is not updated.

## NOTE (Programming)

On some implementations, operations that result in overflow or underflow are likely to take significantly longer than those that do not. For example, these operations may cause a system error handler to be invoked; on such implementations, the system error handler updates overflow bits appropriately.

### 4.4.1.3 Embedded Floating-Point Invalid Operation/Input Errors

Embedded floating-point invalid operation/input errors occur when an operand to an operation contains an invalid input value. If any of the input values are infinity, denorm, or NaN , or for an embedded floating-point divide instruction both operands are $+/-0$, SPEFSCR[FINV FINVH] are set appropriately,
and SPEFSCR[FGH FXH FG FX] are cleared appropriately. If SPEFSCR[FINVE]=1, an embedded floating-point data interrupt is taken and the destination register is not updated.

### 4.4.1.4 Embedded Floating-Point Round (Inexact)

If any result element of an embedded floating-point instruction is inexact, or overflows but embedded floating-point overflow exceptions are disabled, or underflows but embedded floating-point underflow exceptions are disabled, and no higher priority interrupt occurs, SPEFSCR[FINXS] is set. If the embedded floating-point round (inexact) exception is enabled, an embedded floating-point round interrupt occurs. In this case, the destination register is updated with the truncated results. SPEFSCR[FGH FXH FG FX] are properly updated to allow rounding to be performed in the interrupt handler.

SPEFSCR[FG FX] (SPEFSCR[FGH FXH]) are cleared if an embedded floating-point data interrupt is taken due to overflow or underflow, or if an embedded floating-point invalid operation/input error is signaled for the low (high) element (regardless of SPEFSCR[FINVE]).

### 4.4.1.5 Embedded Floating-Point Divide by Zero

If an embedded floating-point divide instruction executes and an embedded floating-point invalid operation/input error does not occur and the instruction is executed with a $+/-0$ divisor value and a finite normalized nonzero dividend value, an embedded floating-point divide by zero exception occurs and SPEFSCR[FDBZ FDBZH] are set appropriately. If embedded floating-point divide by zero exceptions are enabled, an embedded floating-point data interrupt is then taken and the destination register is not updated.

### 4.4.1.6 Default Results

Default results are generated when an embedded floating-point invalid operation/input error, embedded floating-point overflow, embedded floating-point underflow, or embedded floating-point divide by zero occurs on an embedded floating-point operation. Default results provide a normalized value as a result of the operation. In general, denormalized results and underflows are cleared and overflows are saturated to the maximum representable number.

Default results for each operation are described in Section 5.3.4, "Embedded Floating-Point Results."

## Chapter 5

Instruction Set
This chapter describes the SPE instructions and the embedded floating-point instructions, which are as follows:

- Single-precision scalar floating-point (SPE FS)
- Single-precision vector floating-point (SPE FV)
- Double-precision scalar floating-point (SPE FD)


### 5.1 Notation

The definitions and notation listed in Table 5-1 are used throughout this chapter in the instruction descriptions.

Table 5-1. Notation Conventions

| Symbol | Meaning |
| :---: | :---: |
| $\mathrm{X}_{\mathrm{p}}$ | Bit p of register/field X |
| $\mathrm{X}_{\text {field }}$ | The bits composing a defined field of $X$. For example, $X_{\text {sign }}, X_{\text {exp }}$, and $X_{\text {frac }}$ represent the sign, exponent, and fractional value of a floating-point number X |
| $\mathrm{X}_{\mathrm{p}: \mathrm{q}}$ | Bits $p$ through $q$ of register/field $X$ |
| $\mathrm{X}_{\mathrm{pq} \text {.... }}$ | Bits $p, q, \ldots$ of register/field $X$ |
| $\neg \mathrm{X}$ | The one's complement of the contents of $X$ |
| Field i | Bits $4 \times$ through $4 \times i+3$ of a register |
| 1 | Describes the concatenation of two values. For example, $010 \\| 111$ is the same as 010111. |
| $x^{n}$ | $x$ raised to the $\mathrm{n}^{\text {th }}$ power |
| ${ }^{n} \mathrm{x}$ | The replication of $x, n$ times (i.e., $x$ concatenated to itself $n-1$ times). ${ }^{n} 0$ and ${ }^{n} 1$ are special cases: ${ }^{n_{0}}$ means a field of $n$ bits with each bit equal to 0 . Thus ${ }^{5} 0$ is equivalent to $0 b 0 \_0000$. ${ }^{n} 1$ means a field of $n$ bits with each bit equal to 1 . Thus ${ }^{5} 1$ is equivalent to $0 b 1 \_1111$. |
| /, //, /II, | A reserved field in an instruction or in a register. |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

### 5.2 Instruction Fields

Table 5-2 describes instruction fields.
Table 5-2. Instruction Field Descriptions

| Field | Description |
| :--- | :--- |
| CRS (11-13) | Used to specify a CR field to be used as a source |
| D (16-31) | Immediate field used to specify a 16-bit signed two's complement integer that is sign-extended to 64 bits |
| LI (6-29) | Immediate field specifying a 24-bit signed two's complement integer that is concatenated on the right <br> with 0b00 and sign-extended to 64 bits |
| LK (31) | LINK bit. Indicates whether the link register (LR) is set. <br> 0 <br> 1 Do not set the LR. <br> Set the LR. The sum of the value 4 and the address of the branch instruction is placed into the LR. |
| OPCD (0-5) | Primary opcode field |
| rA (11-15) | Used to specify a GPR to be used as a source or as a target |
| rB (16-20) | Used to specify a GPR to be used as a source |
| RS (6-10) | Used to specify a GPR to be used as a source |
| RD (6-10) | Used to specify a GPR to be used as a target |
| SIMM (16-31) | Immediate field used to specify a 16-bit signed integer |
| UIMM (16-31) | Immediate field used to specify a 16-bit unsigned integer |

### 5.3 Description of Instruction Operations

The operation of most instructions is described by a series of statements using a semiformal language at the register transfer level (RTL), which uses the general notation given in Table 5-1 and Table 5-2 and the RTL-specific conventions in Table 5-3. See the example in Figure 5-1. Some of this notation is used in the formal descriptions of instructions.

The RTL descriptions cover the normal execution of the instruction, except that 'standard' setting of the condition register, integer exception register, and floating-point status and control register are not always shown. (Non-standard setting of these registers, such as the setting of condition register field 0 by the stwex. instruction, is shown.) The RTL descriptions do not cover all cases in which exceptions may occur, or for which the results are boundedly undefined, and may not cover all invalid forms.

RTL descriptions specify the architectural transformation performed by the execution of an instruction. They do not imply any particular implementation.

Table 5-3. RTL Notation

| Notation | Meaning |
| :---: | :--- |
| $\leftarrow$ | Assignment |
| $\leftarrow_{f}$ | Assignment in which the data may be reformatted in the target location |
| $\neg$ | NOT logical operator (one's complement) |

Table 5-3. RTL Notation (continued)

| Notation | Meaning |
| :---: | :---: |
| + | Two's complement addition |
| - | Two's complement subtraction, unary minus |
| $\times$ | Multiplication |
| $\div$ | Division (yielding quotient) |
| $+_{\text {dp }}$ | Floating-point addition, double precision |
| -dp | Floating-point subtraction, double precision |
| $\times_{\text {dp }}$ | Floating-point multiplication, double precision |
| ${ }_{\text {¢ }}$ dp | Floating-point division quotient, double precision |
| $+_{\text {sp }}$ | Floating-point addition, single precision |
| -sp | Floating-point subtraction, single precision |
| ${ }_{\text {sf }}$ | Signed fractional multiplication. Result of multiplying two quantities of bit lengths $x$ and $y$ taking the least significant $\mathrm{x}+\mathrm{y}-1$ bits of the product and concatenating a 0 to the Isb forming a signed fractional result of $\mathrm{x}+\mathrm{y}$ bits. |
| $\mathrm{x}_{\text {si }}$ | Signed integer multiplication |
| ${ }_{\text {sp }}$ | Floating-point multiplication, single precision |
| ${ }_{\text {sp }}$ | Floating-point division, single precision |
| $x_{\text {fp }}$ | Floating-point multiplication to infinite precision (no rounding) |
| $\times_{u i}$ | Unsigned integer multiplication |
| =, \# | Equals, Not Equals relations |
| <, $\leq,>, \geq$ | Signed comparison relations |
| $<_{u},>_{u}$ | Unsigned comparison relations |
| ? | Unordered comparison relation |
| \&, 1 | AND, OR logical operators |
| $\oplus$, $\equiv$ | Exclusive OR, Equivalence logical operators ( $(\mathrm{a}=\mathrm{b})=(\mathrm{a} \oplus \rightarrow \mathrm{b})$ ) |
| >>, << | Shift right or left logical |
| ABS(x) | Absolute value of $x$ |
| EXTS(x) | Result of extending x on the left with signed bits |
| EXTZ(x) | Result of extending x on the left with zeros |
| $\operatorname{GPR}(\mathrm{x})$ | General-purpose register x |
| $\operatorname{MASK}(\mathrm{x}, \mathrm{y})$ | Mask having 1s in bit positions x through y (wrapping if $\mathrm{x}>\mathrm{y}$ ) and Os elsewhere |
| MEM (x, 1 ) | Contents of the byte of memory located at address x |
| $\begin{gathered} \text { MEM }(x, y) \\ \text { (for } y=\{2,4,8\}) \end{gathered}$ | Contents of $y$ bytes of memory starting at address $x$. If big-endian memory, the byte at address $x$ is the MSB and the byte at address $x+y-1$ is the LSB of the value being accessed. If little-endian memory, the byte at address $x$ is the LSB and the byte at address $x+y-1$ is the MSB. |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table 5-3. RTL Notation (continued)

| Notation | Meaning |
| :---: | :--- |
| undefined | An undefined value. The value may vary between implementations and between different executions on the <br> same implementation. |
| if ... then $\ldots$ <br> else ... | Conditional execution, indenting shows range; else is optional <br> doDo loop, indenting shows range. 'To' and/or 'by' clauses specify incrementing an iteration variable, and a 'while' <br> clause gives termination conditions. |

Precedence rules for RTL operators are summarized in Table 5-4. Operators higher in the table 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. (For example, - associates from left to right, so $\mathrm{a}-\mathrm{b}-\mathrm{c}=(\mathrm{a}-\mathrm{b})-\mathrm{c}$.) Parentheses are used to override the evaluation order implied by the table or to increase clarity; parenthesized expressions are evaluated before serving as operands.

Table 5-4. Operator Precedence

| Operators | Associativity |
| :--- | :---: |
| Subscript, function evaluation | Left to right |
| Pre-superscript (replication), post-superscript (exponentiation) | Right to left |
| unary,$- \neg$ | Right to left |
| $\times, \div$ | Left to right |
| ,+- | Left to right |
| II | Left to right |
| $=, \neq,<, \leq,>, \geq,<_{u},>_{u}, ?$ | Left to right |
| $\&, \oplus, \equiv$ | Left to right |
| I | Left to right |
| $:($ range $)$ | None |
| $\leftarrow$ | None |

### 5.3.1 SPE Saturation and Bit-Reverse Models

For saturation and bit reversal, the pseudo RTL is provided here to more accurately describe those functions that are referenced in the instruction pseudo RTL.

### 5.3.1.1 Saturation

```
SATURATE(overflow, carry, saturated_underflow, saturated_overflow, value)
if overflow then
    if carry then
        return saturated_underflow
    else
        return saturated_overflow
```

```
else
```

    return value
    
### 5.3.1.2 Bit Reverse

BITREVERSE (value)

```
result \leftarrow 0
mask \leftarrow }
shift \leftarrow }\leftarrow3
cnt \leftarrow }\leftarrow3
while cnt > 0 then do
    t data & mask
    if shift >= 0 then
        result \leftarrow(t << shift) result
    else
        result \leftarrow(t >> -shift) | result
    cnt \leftarrow cnt - 1
    shift \leftarrowshift - 2
    mask \leftarrowmask << 1
return result
```


### 5.3.2 Embedded Floating-Point Conversion Models

The embedded floating-point instructions defined by the signal processing engine (SPE) contain floating-point conversion to and from integer and fractional type instructions. The floating-point to-and-from non-floating-point conversion model pseudo-RTL is provided in Table 5-5 as a group of functions that is called from the individual instruction pseudo-RTL descriptions.

Table 5-5. Conversion Models

| Function | Name | Reference |
| :--- | :--- | :--- |
| Common Functions |  |  |
| Round a 32-bit value | Round32(fp,guard,sticky) | $5.3 .2 .1 .3 / 5-6$ |
| Round a 64-bit value | Round64(fp,guard,sticky) | $5.3 .2 .1 .4 / 5-7$ |
| Signal floating-point error | SignalFPError | $5.3 .2 .1 .2 / 5-6$ |
| Is a 32-bit value a NaN or infinity? | Isa32NaNorinfinity(fp) | $5.3 .2 .1 .1 / 5-6$ |
| Floating-Point Conversions |  |  |
| Convert from single-precision floating-point to <br> integer word with saturation | CnvtFP32Tol32Sat(fp,signed,upper_lower,round,fractional) | $5.3 .2 .2 / 5-7$ |
| Convert from double-precision floating-point to <br> integer word with saturation | CnvtFP64Tol32Sat(fp,signed,round,fractional) | $5.3 .2 .3 / 5-9$ |
| Convert from double-precision floating-point to <br> integer double word with saturation | CnvtFP64Tol64Sat(fp,signed,round) | $5.3 .2 .4 / 5-10$ |
| Convert to single-precision floating-point from <br> integer word with saturation | Cnvtl32ToFP32Sat(v,signed,upper_lower,fractional) | $5.3 .2 .5 / 5-11$ |
| Convert to double-precision floating-point from <br> integer double word with saturation | Cnvtl64ToFP64Sat(v,signed) | $5.3 .2 .7 / 5-13$ |

Table 5-5. Conversion Models (continued)

| Function | Name | Reference |  |  |
| :--- | :--- | :---: | :---: | :---: |
| Integer Saturate |  |  |  | $5.3 .3 / 5-14$ |
| Integer saturate | SATURATE(ovf,carry,neg_sat,pos_sat,value) |  |  |  |

### 5.3.2.1 Common Embedded Floating-Point Functions

This section includes common functions used by the functions in subsequent sections.

### 5.3.2.1.1 $\quad 32$-Bit NaN or Infinity Test

```
// Determine if fp value is a NaN or infinity
Isa32NaNorInfinity(fp)
return (fpexp = 255)
Isa32NaN(fp)
return ((fpexp = 255) & (fppfrac fo 0))
Isa32Infinity(fp)
return ((fpexp = 255) & (fppfrac = 0))
// Determine if fp value is denormalized
Isa32Denorm(fp)
return ((f\mp@subsup{p}{\operatorname{exp}}{}=0) & (f\mp@subsup{p}{\mathrm{ frac }}{\prime}=0))
// Determine if fp value is a NaN or Infinity
Isa64NaNorInfinity(fp)
return (fpexp = 2047)
Isa64NaN(fp)
return ((fpexp = 2047) & (f\mp@subsup{p}{\mathrm{ frac }}{\prime=0))}
Isa64Infinity(fp)
return ((f\mp@subsup{p}{exp}{}=2047) & (f\mp@subsup{p}{\mathrm{ frac }}{\prime}=0))
// Determine if fp value is denormalized
Isa64Denorm(fp)
return ((f\mp@subsup{p}{exp}{e}=0) & (f\mp@subsup{p}{frac}{f}\not=0))
```


### 5.3.2.1.2 Signal Floating-Point Error

```
// Signal a Floating-Point Error in the SPEFSCR
SignalFPError(upper_lower, bits)
if (upper_lower = UPPER) then
    bits }\leftarrow bits << 1
SPEFSCR \leftarrow SPEFSCR | bits
bits \leftarrow (FG | FX)
if (upper_lower = UPPER) then
    bits }\leftarrow bits << 1
SPEFSCR \leftarrow SPEFSCR & 子bits
```


### 5.3.2.1.3 Round a 32-Bit Value

```
// Round a result
Round32(fp, guard, sticky)
FP32format fp;
if (SPEFSCR FINXE = 0) then
    if (SPEFSCR FRMC = 0b00) then // nearest
        if (guard) then
            if (sticky | fpfrac[22]) then
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

```
v[0:23]}\leftarrow f\mp@subsup{p}{\mathrm{ frac }}{}+
if v[0] then
    if (fppexp >= 254) then
        // overflow
        fp}\leftarrow\mp@subsup{\textrm{fp}}{\mathrm{ sign }}{}||0\textrm{b}11111110||\mp@subsup{}{}{23}
    else
        fp
        fp_frac}\leftarrow\leftarrow\mp@subsup{v}{1:23}{
else
    fp}\mp@subsup{f}{\mathrm{ frac }}{}\leftarrowv[1:23
    else if ((SPEFSCR FRMC & 0b10) = 0b10) then // infinity modes
    // implementation dependent
return fp
```


### 5.3.2.1.4 Round a 64-Bit Value

```
// Round a result
Round64(fp, guard, sticky)
FP32format fp;
if (SPEFSCR FINXE = 0) then
    if (SPEFSCR FRMC = 0b00) then // nearest
        if (guard) then
            if (sticky | fpfrac[51]) then
                v[0:52]}\leftarrow\mp@subsup{\textrm{fp}}{\mathrm{ frac }}{}+
                if v[0] then
                    if (fpexp >= 2046) then
                        // overflow
                        fp}\leftarrow\mp@subsup{\textrm{fp}}{\mathrm{ sign }}{}||0b11111111110 || 521
                    else
                        fp
                        fp}\mp@subsup{f}{\mathrm{ frac }}{}\leftarrow\mp@subsup{\textrm{V}}{1:52}{
            else
            fp
    else if ((SPEFSCR FRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
return fp
```


### 5.3.2.2 Convert from Single-Precision Floating-Point to Integer Word with Saturation

```
// Convert 32-bit floating point to integer/factional
// signed = SIGN or UNSIGN
// upper_lower = UPPER or LOWER
// round = ROUND or TRUNC
// fractional = F (fractional) or I (integer)
CnvtFP32ToI32Sat(fp, signed, upper_lower, round, fractional)
FP32format fp;
if (Isa32NaNorInfinity(fp)) then // SNaN, QNaN, +-INF
    SignalFPError(upper_lower, FINV)
    if (Isa32NaN(fp)) then
    return 0x00000000 // all NaNs
    if (signed = SIGN) then
        if (fp
            return 0x80000000
        else
            return 0x7fffffff
    else
        if (fp
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Instruction Set

```
            return 0x00000000
        else
            return Oxffffffff
if (Isa32Denorm(fp)) then
        SignalFPError(upper_lower, FINV)
        return 0x00000000 // regardless of sign
if ((signed = UNSIGN) & (fp
        SignalFPError(upper_lower, FOVF) // overflow
        return 0x00000000
if ((fpexp = 0) & (fpprrac = 0)) then
    return 0x00000000 // all zero values
if (fractional = I) then // convert to integer
    max_exp \leftarrow 158
    shift \leftarrow 158 - fpeexp
    if (signed = SIGN) then
        if ((f\mp@subsup{p}{\operatorname{exp}}{}\not=158) | (f\mp@subsup{p}{\textrm{frac}}{\prime}\not=0) | (f\mp@subsup{p}{\mathrm{ sign }}{\prime}\not=1)) then
            max_exp \leftarrow max_exp - 1
else // fractional conversion
    max_exp \leftarrow 126
    shift}\leftarrow126 - fpexp
    if (signed = SIGN) then
        shift \leftarrow shift + 1
if (fpexp > max_exp) then
    SignalFPError(upper_lower, FOVF) // overflow
    if (signed = SIGN) then
        if (fppsign = 1) then
            return 0x80000000
        else
            return 0x7fffffff
    else
        return 0xfffffffff
result \leftarrow0bl || fp frac || 0b00000000 // add U to frac
guard \leftarrow0
sticky \leftarrow 0
for (n \leftarrow 0; n < shift; n \leftarrow n + 1) do
    sticky \leftarrow sticky | guard
    guard \leftarrow result & 0x00000001
    result \leftarrow result > 1
// Report sticky and guard bits
if (upper_lower = UPPER) then
    SPEFSCR FGH }\leftarrow\mathrm{ guard
    SPEFSCR FXH}\leftarrow\leftarrow sticky
else
    SPEFSCR FG
    SPEFSCR FX }\leftarrow stick
if (guard | sticky) then
    SPEFSCR FINXS }\leftarrow
// Round the integer result
if ((round = ROUND) & (SPEFSCR FINXE = 0)) then
    if (SPEFSCR FRMC = 0b00) then // nearest
        if (guard) then
                if (sticky | (result & 0x00000001)) then
                    result \leftarrow result + 1
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

```
    else if ((SPEFSCR FRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
if (signed = SIGN) then
    if (fppsign = 1) then
        result \leftarrow \negresult + 1
return result
```


### 5.3.2.3 Convert from Double-Precision Floating-Point to Integer Word with Saturation

```
// Convert 64-bit floating point to integer/fractional
// signed = SIGN or UNSIGN
// round = ROUND or TRUNC
// fractional = F (fractional) or I (integer)
CnvtFP64ToI32Sat(fp, signed, round, fractional)
FP64format fp;
if (Isa64NaNorInfinity(fp)) then // SNaN, QNaN, +-INF
    SignalFPError(LOWER, FINV)
    if (Isa64NaN(fp)) then
        return 0x00000000 // all NaNs
    if (signed = SIGN) then
        if (fp
            return 0x800000000
        else
            return 0x7fffffff
    else
        if (fp
            return 0x000000000
        else
            return 0xffffffff
if (Isa64Denorm(fp)) then
    SignalFPError(LOWER, FINV)
    return 0x00000000 // regardless of sign
if ((signed = UNSIGN) & (fppsign = 1)) then
    SignalFPError(LOWER, FOVF) // overflow
    return 0x000000000
if ((fpexp =0) & (fp
    return 0x00000000 // all zero values
if (fractional = I) then // convert to integer
    max_exp \leftarrow 1054
    shift \leftarrow 1054 - fpexp
    if (signed \leftarrow SIGN) then
        if ((fpexp f 1054) | (f\mp@subsup{p}{\mathrm{ frac }}{\prime\not=0) | (fp sign fl 1)) then}
            max_exp \leftarrow max_exp - 1
else - // fractional conversion
    max_exp \leftarrow 1022
    shift \leftarrow 1022 - fpexp
    if (signed = SIGN) then
        shift \leftarrow shift + 1
if (fpexp > max_exp) then
    SignalFPError(LOWER, FOVF) // overflow
    if (signed = SIGN) then
        if (fp
            return 0x80000000
        else
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Instruction Set

```
                return 0x7fffffff
    else
    return Oxffffffff
result }\leftarrow0.b1 || fp frac [0:30] // add U to fra
guard }\leftarrow\mp@subsup{\textrm{fp}}{\mathrm{ frac[31]}}{
sticky }\leftarrow(\mp@subsup{f}{\mathrm{ frac [32:63] }}{=00)
for (n \leftarrow 0; n < shift; n \leftarrow n + 1) do
    sticky \leftarrow sticky | guard
    guard }\leftarrow\mathrm{ result & 0x00000001
    result \leftarrow result > 1
// Report sticky and guard bits
SPEFSCR 
SPEFSCR 
if (guard | sticky) then
    SPEFSCR FINXS 
// Round the result
if ((round = ROUND) & (SPEFSCR FINXE = 0)) then
    if (SPEFSCR FRMC = 0.b00) then // nearest
        if (guard) then
                if (sticky | (result & 0x00000001)) then
                        result }\leftarrow result + 1
    else if ((SPEFSCR FRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
if (signed = SIGN) then
    if (fppign = 1) then
        result }\leftarrow \negresult + 1
return result
```


### 5.3.2.4 Convert from Double-Precision Floating-Point to Integer Double Word with Saturation

```
// Convert 64-bit floating point to integer/fractional
// signed = SIGN or UNSIGN
// round = ROUND or TRUNC
CnvtFP64ToI64Sat(fp, signed, round)
FP64format fp;
if (Isa64NaNorInfinity(fp)) then // SNaN, QNaN, +-INF
    SignalFPError(LOWER, FINV)
    if (Isa64NaN(fp)) then
        return 0x00000000_00000000 // all NaNs
    if (signed = SIGN) then
        if (fp
            return 0x80000000_000000000
        else
            return 0x7ffffffff_ffffffff
    else
        if (fp
            return 0x00000000_000000000
            else
            return 0xfffffffff_fffffffff
if (Isa64Denorm(fp)) then
    SignalFPError(LOWER, FINV)
    return 0x00000000_00000000 // regardless of sign
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

```
if ((signed = UNSIGN) & (fppsign = 1)) then
    SignalFPError(LOWER, FOVF) // overflow
    return 0x00000000_00000000
if ((fpexp =0) & (fppfrac =0)) then
    return 0x00000000_00000000 // all zero values
max_exp \leftarrow 1086
shift }\leftarrow1086 - fpexp
if (signed = SIGN) then
    if ((f\mp@subsup{p}{\operatorname{exp}}{}\not=1086) | (f\mp@subsup{p}{\mathrm{ frac }}{}\not=0) | (f\mp@subsup{p}{\mathrm{ sign }}{\prime= 1)) then}
        max_exp \leftarrow max_exp - 1
if (fpexp > max_exp) then
    SignalFPError(LOWER, FOVF) // overflow
    if (signed = SIGN) then
        if (fppsign = 1) then
            return 0x80000000_000000000
            else
            return 0x7fffffff_ffffffff
    else
        return 0xffffffff_ffffffff
result \leftarrow 0bl || fpfrac || 0b00000000000 // add U to frac
guard \leftarrow0
sticky }\leftarrow
for (n \leftarrow 0; n < shift; n \leftarrow n + 1) do
    sticky \leftarrow sticky | guard
    guard \leftarrow result & 0x00000000_000000001
    result \leftarrow result > 1
// Report sticky and guard bits
SPEFSCR FG
SPEFSCR FX }\leftarrow stick
if (guard | sticky) then
    SPEFSCR FINXS }\leftarrow
// Round the result
if ((round = ROUND) & (SPEFSCR FINXE = 0)) then
    if (SPEFSCR FRMC = 0b00) then // nearest
        if (guard) then
            if (sticky | (result & 0x00000000_00000001)) then
                result \leftarrow result + 1
    else if ((SPEFSCR FRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
if (signed = SIGN) then
    if (fppsign = 1) then
        result \leftarrow \negresult + 1
return result
```


### 5.3.2.5 Convert to Single-Precision Floating-Point from Integer Word with Saturation

```
// Convert from integer/factional to 32-bit floating point
// signed = SIGN or UNSIGN
// upper_lower = UPPER or LOWER
// fractional = F (fractional) or I (integer)
CnvtI32ToFP32Sat(v, signed, upper_lower, fractional)
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Instruction Set

```
FP32format result;
result sign }\leftarrow
if (v = 0) then
    result \leftarrow 0
    if (upper_lower = UPPER) then
        SPEFSCR FGH }\leftarrow
        SPEFSCR FXH
        else
            SPEFSCR FG
            SPEFSCR FX }\leftarrow
else
    if (signed = SIGN) then
            if ( }\mp@subsup{v}{0}{\prime}=1) the
                v \leftarrow ᄀv + 1
                result sign }\leftarrow
        if (fractional = F) then // fractional bit pos alignment
        maxexp \leftarrow < 127
        if (signed = UNSIGN) then
            maxexp }\leftarrow\mathrm{ maxexp - 1
        else
            maxexp \leftarrow 158 // integer bit pos alignment
        sc}\leftarrow
        while (vo = 0)
            v \leftarrow v << 1
            sc}\leftarrow\textrm{sc}+
        vo \leftarrow 0 // clear U bit
        result exp }\leftarrow\mathrm{ maxexp - sc
        guard }\leftarrow\mp@subsup{\textrm{v}}{24}{
        sticky }\leftarrow(\mp@subsup{v}{25:31}{}\not=0
        // Report sticky and guard bits
        if (upper_lower = UPPER) then
            SPEFSCR FGG }\leftarrow\mathrm{ guard
            SPEFSCR 
        else
            SPEFSCR 
            SPEFSCR FX }\leftarrow stick
        if (guard | sticky) then
            SPEFSCR FINXS
// Round the result
        result frac }\leftarrow\mp@subsup{\textrm{v}}{1:23}{
        result \leftarrow Round32(result, guard, sticky)
    return result
```


### 5.3.2.6 Convert to Double-Precision Floating-Point from Integer Word with Saturation

```
// Convert from integer/factional to 64-bit floating point
// signed = SIGN or UNSIGN
// fractional = F (fractional) or I (integer)
CnvtI32ToFP64Sat(v, signed, fractional)
FP64format result;
result sign }\leftarrow
if (v = 0) then
    result }\leftarrow
    SPEFSCR FG}\leftarrow
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

```
    SPEFSCR RX 
else
    if (signed = SIGN) then
        if (v[0] = 1) then
                v \leftarrow ᄀv + 1
                result sign }\leftarrow
    if (fractional = F) then // fractional bit pos alignment
        maxexp \leftarrow 1023
        if (signed = UNSIGN) then
            maxexp \leftarrow maxexp - 1
    else
        maxexp \leftarrow 1054 // integer bit pos alignment
    sc}\leftarrow
    while (vo = 0)
        v \leftarrow v << 1
        sc}\leftarrow sc + 1
    vo \leftarrow 0 // clear U bit
    result exp }\leftarrow maxexp - s
// Report sticky and guard bits
    SPEFSCR 
    SPEFSCR FX }\leftarrow
    result frac }\leftarrow\mp@subsup{v}{1:31}{}||\mp@subsup{}{}{21}
return result
```


### 5.3.2.7 Convert to Double-Precision Floating-Point from Integer Double Word with Saturation

```
// Convert from 64 integer to 64-bit floating point
// signed = SIGN or UNSIGN
CnvtI64ToFP64Sat(v, signed)
FP64format result;
result sign }\leftarrow
if (v = 0) then
    result }\leftarrow
    SPEFSCR 
    SPEFSCR FX }\leftarrow
else
    if (signed = SIGN) then
        if ( }\mp@subsup{v}{0}{}=1\mathrm{ ) then
            v \leftarrow ᄀv + 1
            result sign }\leftarrow
    maxexp \leftarrow }\leftarrow105
    sc}\leftarrow
    while (vo = 0)
        v \leftarrow v << 1
        sc}\leftarrow\textrm{sc}+
    vo \leftarrow 0 // clear U bit
    result exp }\leftarrow\mathrm{ maxexp - sc
    guard }\leftarrow\mp@subsup{v}{53}{
    sticky \leftarrow (\mp@subsup{v}{54:63 }{*}=0)
// Report sticky and guard bits
    SPEFSCR 
    SPEFSCR FX }\leftarrow sticky
    if (guard | sticky) then
        SPEFSCR FINXS 
```

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

```
// Round the result
    result frac }\leftarrow\mp@subsup{\textrm{v}}{1:52}{
    result }\leftarrow\mathrm{ Round64(result, guard, sticky)
return result
```


### 5.3.3 Integer Saturation Models

```
// Saturate after addition
SATURATE(ovf, carry, neg_sat, pos_sat, value)
if ovf then
    if carry then
        return neg_sat
    else
        return pos_sat
else
    return value
```


### 5.3.4 Embedded Floating-Point Results

Section 5.3.4, "Embedded Floating-Point Results," summarizes results of various types of SPE and embedded floating-point operations on various combinations of input operands.

### 5.4 Instruction Set

The rest of this chapter describes individual instructions, which are listed in alphabetical order by mnemonic. Figure 5-1 shows the format for instruction description pages.


Figure 5-1. Instruction Description
Note that the execution unit that executes the instruction may not be the same for all processors.

## brinc

| SPE | User |
| :---: | :---: |

Bit Reversed Increment
brinc rD,rA,rB


```
n < MASKBITS // Imp dependent # of mask bits
mask \leftarrow rB64-n:63 // Least sig. n bits of register
a}\leftarrowr\mp@subsup{r}{64-n:63}{
d \leftarrow bitreverse(1 + bitreverse(a | (\negmask)))
rD}\leftarrowr\mp@subsup{A}{0:63-n}{}||(d& mask
```

brinc provides a way for software to access FFT data in a bit-reversed manner. rA contains the index into a buffer that contains data on which FFT is to be performed. $\mathbf{r B}$ contains a mask that allows the index to be updated with bit-reversed addressing. Typically this instruction precedes a load with index instruction; for example,

```
brinc r2, r3, r4
lhax r8, r5, r2
```

$\mathbf{r B}$ contains a bit-mask that is based on the number of points in an FFT. To access a buffer containing n byte sized data that is to be accessed with bit-reversed addressing, the mask has $\log _{2} n 1 \mathrm{~s}$ in the least significant bit positions and 0 s in the remaining most significant bit positions. If, however, the data size is a multiple of a half word or a word, the mask is constructed so that the 1 s are shifted left by $\log _{2}$ (size of the data) and 0 s are placed in the least significant bit positions. Table 5-6 shows example values of masks for different data sizes and number of data.

Table 5-6. Data Samples and Sizes

| Number of Data Samples | Data Size |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  | Byte | Half Word | Word | Double Word |
| 8 | $000 \ldots 00000111$ | $000 \ldots 00001110$ | $000 \ldots 000011100$ | $000 \ldots 0000111000$ |
| 16 | $000 \ldots 00001111$ | $000 \ldots 00011110$ | $000 \ldots 000111100$ | $000 \ldots 0001111000$ |
| 32 | $000 \ldots 00011111$ | $000 \ldots 00111110$ | $000 \ldots 001111100$ | $000 \ldots 0011111000$ |
| 64 | $000 \ldots 00111111$ | $000 \ldots 01111110$ | $000 \ldots 011111100$ | $000 \ldots 0111111000$ |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## efdabs

SPE FD $\quad$ User

## efdabs

Floating-Point Double-Precision Absolute Value
efdabs rD,rA

$r D_{0: 63} \leftarrow 0 \mathrm{bo} \| r \mathrm{~A}_{1: 63}$
The sign bit of $\mathbf{r A}$ is set to 0 and the result is placed into $\mathbf{r D}$.
Exceptions:
Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the sign bit operation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: If $\mathbf{r A}$ is infinity, denorm, or NaN, SPEFSCR[FINV] is set, and FG and FX are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the destination register is not updated.

## efdadd

SPE FD $\quad$ User

Floating-Point Double-Precision Add
efdadd rD,rA,rB


$$
r D_{0: 63} \leftarrow \mathrm{rA}_{0: 63}+{ }_{\mathrm{dp}} \mathrm{rB}_{0: 63}
$$

$\mathbf{r A}$ is added to $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$. If $\mathbf{r A}$ is NaN or infinity, the result is either $\operatorname{pmax}\left(\mathrm{a}_{\text {sign }}==0\right)$, or $n \max \left(\mathrm{a}_{\text {sign }}==1\right)$. Otherwise, If $\mathbf{r B}$ is NaN or infinity, the result is either pmax $\left(\mathrm{b}_{\text {sign }}==0\right)$, or $n$ max $\left(b_{\text {sign }}==1\right)$. Otherwise, if an overflow occurs, pmax or $n \max$ (as appropriate) is stored in $\mathbf{r D}$. If an underflow occurs, +0 (for rounding modes RN, RZ, RP) or -0 (for rounding mode RM) is stored in rD.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF] is set, or if an underflow occurs, SPEFSCR[FUNF] is set. If either underflow or overflow exceptions are enabled and the corresponding bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If the result of this instruction is inexact or if an overflow occurs but overflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
FG and FX are cleared if an overflow, underflow, or invalid operation/input error is signaled, regardless of enabled exceptions.

## efdcfs

| SP.FD | User |
| :--- | :--- |

Floating-Point Double-Precision Convert from Single-Precision
efdcfs rD,rB


```
FP32format f;
FP64format result;
\(\mathrm{f} \leftarrow \mathrm{rB}_{32: 63}\)
if \(\left.\left(f_{\text {exp }}=0\right) \&\left(f_{\text {frac }}=0\right)\right)\) then
    result \(\leftarrow \mathrm{f}_{\text {sign }}| |{ }^{63} 0 \quad / /\) signed zero value
else if Isa32NaNorInfinity(f) | Isa32Denorm(f) then
    SPEFSCR \(_{\text {FINV }} \leftarrow 1\)
    result \(\leftarrow \mathrm{f}_{\text {sign }}| | 0\) b11111111110 || \({ }^{52} 1\) // max value
else if Isa32Denorm(f) then
    \(\operatorname{SPEFSCR}_{\text {FINV }} \leftarrow 1\)
    result \(\leftarrow \mathrm{f}_{\text {sign }}| |{ }^{63} 0\)
else
    result \(_{\text {sign }} \leftarrow \mathrm{f}_{\text {sign }}\)
    result \({ }_{\text {exp }} \leftarrow \mathrm{f}_{\text {exp }}-1277^{+1023}\)
result
\(r D_{0: 63}=\) result
```

The single-precision floating-point value in the low element of $\mathbf{r B}$ is converted to a double-precision floating-point value and the result is placed into $\mathbf{r D}$. The rounding mode is not used since this conversion is always exact.

## Exceptions:

If the low element of $\mathbf{r B}$ is infinity, denorm, or NaN, SPEFSCR[FINV] is set. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.
FG and FX are always cleared.

## efdcfsf

## Convert Floating-Point Double-Precision from Signed Fraction

efdcfsf rD,rB

$r D_{0: 63} \leftarrow$ CnvtI32ToFP64 (rB ${ }_{32: 63}$, SIGN, $\left.F\right)$
The signed fractional low element in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into $\mathbf{r D}$.

Exceptions:
None.

## efdcfsi

SPE FD $\quad$ User

## Convert Floating-Point Double-Precision from Signed Integer

efdcfsi rD,rB

$r D_{0: 63} \leftarrow$ CnvtSI32ToFP64 (rB $32: 63$, SIGN, I)
The signed integer low element in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into rD.

Exceptions:
None.

## efdcfsid

SPE FD $\quad$ User

## Convert Floating-Point Double-Precision from Signed Integer Doubleword

efdcfsid rD,rB


The signed integer doubleword in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into rD.

## Exceptions:

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

This instruction may only be implemented for 64-bit implementations.

## efdcfuf

## Convert Floating-Point Double-Precision from Unsigned Fraction

efdcfuf rD,rB

$r D_{0: 63} \leftarrow$ CnvtI32ToFP64 (rB 32:63 , UNSIGN, $F$ )
The unsigned fractional low element in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into rD.

Exceptions:
None.

## efdcfui

## Convert Floating-Point Double-Precision from Unsigned Integer

efdcfui rD,rB


The unsigned integer low element in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into $\mathbf{r D}$.

Exceptions:
None.

## efdcfuid

SPE FD $\quad$ User

## efdcfuid

Convert Floating-Point Double-Precision from Unsigned Integer Doubleword
efdcfuid rD,rB


The unsigned integer doubleword in $\mathbf{r B}$ is converted to a double-precision floating-point value using the current rounding mode and the result is placed into rD.

## Exceptions:

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

This instruction may only be implemented for 64-bit implementations.

## efdcmpeq

## Floating-Point Double-Precision Compare Equal

efdcmpeq crfD,rA,rB

$\mathrm{al} \leftarrow r \mathrm{~A}_{0: 63}$
$\mathrm{bl} \leftarrow r \mathrm{~B}_{0: 63}$
if (al = bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4 * \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is equal to $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared.
Comparison ignores the sign of $0(+0=-0)$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set, and the FGH FXH, FG and FX bits are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## efdcmpgt

SPE FD $\quad$ User

Floating-Point Double-Precision Compare Greater Than
efdempgt crfD,rA,rB

al $\leftarrow r A_{0: 63}$
bl $\leftarrow \mathrm{rB}_{0: 63}$
if (al > bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \text { crD }}$ 4* $^{\text {crD }+3} \leftarrow$ undefined || Cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is greater than $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set, and the FGH FXH, FG and FX bits are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## efdcmplt

Floating-Point Double-Precision Compare Less Than
efdemplt crfD,rA,rB

$\mathrm{al} \leftarrow r \mathrm{~A}_{0: 63}$
bl $\leftarrow \mathrm{rB}_{0: 63}$
if (al < bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \text { CrD }: 4 * \mathrm{CrD}+3} \leftarrow$ undefined || Cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

## Exceptions:

If the contents of $\mathbf{r} A$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set, and the FGH FXH, FG and FX bits are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## efdctsf

SPE FD $\quad$ User

## Convert Floating-Point Double-Precision to Signed Fraction

efdctsf rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to a signed fraction using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit fraction. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efdctsi

## Convert Floating-Point Double-Precision to Signed Integer

efdctsi rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to a signed integer using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efdctsidz

## Convert Floating-Point Double-Precision to Signed Integer Doubleword with Round toward Zero

efdctsidz rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to a signed integer doubleword using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 64-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

This instruction may only be implemented for 64-bit implementations.

## efdctsiz

Convert Floating-Point Double-Precision to Signed Integer with Round toward Zero
efdctsiz rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to a signed integer using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.
This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efdctuf

SPE FD $\quad$ User

## Convert Floating-Point Double-Precision to Unsigned Fraction

efdctuf rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned fraction using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit unsigned fraction. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the Floating-Point Round Interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efdctui

SPE FD $\quad$ User

Convert Floating-Point Double-Precision to Unsigned Integer
efdctui rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned integer using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efdctuidz

SPE FD $\quad$ User

## efdctuidz

Convert Floating-Point Double-Precision to Unsigned Integer Doubleword with Round toward Zero
efdctuidz rD,rB

$r D_{0: 63} \leftarrow$ CnvtFP64ToI64Sat (rB $0_{0: 63}$, UNSIGN, TRUNC)
The double-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned integer doubleword using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 64 -bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
This instruction may only be implemented for 64-bit implementations.

## efdctuiz

SPE FD $\quad$ User

## efdctuiz

Convert Floating-Point Double-Precision to Unsigned Integer with Round toward Zero
efdctuiz rD,rB


The double-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned integer using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 32 -bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of $\mathbf{r B}$ are infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV] is set, and the FG, and FX bits are cleared. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversion is not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## efddiv

SPE FD $\quad$ User

Floating-Point Double-Precision Divide
efddiv rD,rA,rB


$$
r \mathrm{D}_{0: 63} \leftarrow \mathrm{rA}_{0: 63} \leftarrow_{\mathrm{dp}} \mathrm{rB}_{0: 63}
$$

$\mathbf{r A}$ is divided by $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$. If $\mathbf{r B}$ is a NaN or infinity, the result is a properly signed zero. Otherwise, if $\mathbf{r B}$ is a zero (or a denormalized number optionally transformed to zero by the implementation), or if $\mathbf{r A}$ is either NaN or infinity, the result is either pmax $\left(\mathrm{a}_{\text {sign }}==\mathrm{b}_{\text {sign }}\right)$, or nmax $\left(a_{\text {sign }}!=b_{\text {sign }}\right)$. Otherwise, if an overflow occurs, pmax or nmax (as appropriate) is stored in $\mathbf{r D}$. If an underflow occurs, +0 or -0 (as appropriate) is stored in $\mathbf{r D}$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN , or if both $\mathbf{r A}$ and $\mathbf{r B}$ are $+/-0$, SPEFSCR[FINV] is set. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated.
Otherwise, if the content of $\mathbf{r B}$ is $+/-0$ and the content of $\mathbf{r} A$ is a finite normalized non-zero number, SPEFSCR[FDBZ] is set. If floating-point divide by zero Exceptions are enabled, an interrupt is then taken. Otherwise, if an overflow occurs, SPEFSCR[FOVF] is set, or if an underflow occurs, SPEFSCR[FUNF] is set. If either underflow or overflow exceptions are enabled and the corresponding bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If the result of this instruction is inexact or if an overflow occurs but overflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

FG and FX are cleared if an overflow, underflow, divide by zero, or invalid operation/input error is signaled, regardless of enabled exceptions.

## efdmul

SPE FD $\quad$ User

## efdmul

Floating-Point Double-Precision Multiply
efdmul rD,rA,rB


$$
r D_{0: 63} \leftarrow \mathrm{rA}_{0: 63} \times_{\mathrm{dp}} \mathrm{rB}_{0: 63}
$$

$\mathbf{r A}$ is multiplied by $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$. If $\mathbf{r A}$ or $\mathbf{r B}$ are zero (or a denormalized number optionally transformed to zero by the implementation), the result is a properly signed zero. Otherwise, if $\mathbf{r A}$ or $\mathbf{r B}$ are either NaN or infinity, the result is either $\operatorname{pmax}\left(a_{\text {sign }}=\mathrm{b}_{\text {sign }}\right)$, or $n \max \left(a_{\text {sign }}!=\mathrm{b}_{\text {sign }}\right)$. Otherwise, if an overflow occurs, pmax or nmax (as appropriate) is stored in rD. If an underflow occurs, +0 or -0 (as appropriate) is stored in $\mathbf{r D}$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF] is set, or if an underflow occurs, SPEFSCR[FUNF] is set. If either underflow or overflow exceptions are enabled and the corresponding bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If the result of this instruction is inexact or if an overflow occurs but overflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
FG and FX are cleared if an overflow, underflow, or invalid operation/input error is signaled, regardless of enabled exceptions.

## efdnabs

SPE FD $\quad$ User

Floating-Point Double-Precision Negative Absolute Value
efdnabs rD,rA

$r D_{0: 63} \leftarrow 0 \mathrm{~b} 1| | r A_{1: 63}$
The sign bit of $\mathbf{r A}$ is set to 1 and the result is placed into $\mathbf{r D}$.
Exceptions:
Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the sign bit operation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: If $\mathbf{r A}$ is infinity, denorm, or NaN, SPEFSCR[FINV] is set, and FG and FX are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the destination register is not updated.

## efdneg

SPE FD $\quad$ User

Floating-Point Double-Precision Negate
efdneg rD,rA


$$
r D_{0: 63} \leftarrow \neg r A_{0} \| r A_{1: 63}
$$

The sign bit of $\mathbf{r A}$ is complemented and the result is placed into $\mathbf{r D}$.
Exceptions:
Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the sign bit operation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: If $\mathbf{r A}$ is infinity, denorm, or NaN, SPEFSCR[FINV] is set, and FG and FX are cleared. If floating-point invalid input exceptions are enabled, an interrupt is taken and the destination register is not updated.

## efdsub

SPE FD $\quad$ User

Floating-Point Double-Precision Subtract
efdsub rD,rA,rB


$$
r D_{0: 63} \leftarrow \mathrm{rA}_{0: 63}-{ }_{\mathrm{dp}} \mathrm{rB}_{0: 63}
$$

$\mathbf{r B}$ is subtracted from rA and the result is stored in $\mathbf{r D}$. If $\mathbf{r A}$ is NaN or infinity, the result is either pmax $\left(a_{\text {sign }}=0\right)$, or $n \max \left(a_{\text {sign }}=1\right)$. Otherwise, If $\mathbf{r B}$ is NaN or infinity, the result is either $n \max \left(\mathrm{~b}_{\text {sign }}==0\right)$, or pmax ( $\mathrm{b}_{\text {sign }}==1$ ). Otherwise, if an overflow occurs, pmax or nmax (as appropriate) is stored in $\mathbf{r D}$. If an underflow occurs, +0 (for rounding modes RN, RZ, RP) or -0 (for rounding mode RM) is stored in $\mathbf{r D}$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV] is set. If SPEFSCR[FINVE] is set, an interrupt is taken, and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF] is set, or if an underflow occurs, SPEFSCR[FUNF] is set. If either underflow or overflow exceptions are enabled and the corresponding bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If the result of this instruction is inexact or if an overflow occurs but overflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result, the FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
FG and FX are cleared if an overflow, underflow, or invalid operation/input error is signaled, regardless of enabled exceptions.

## efdtsteq

| SPE FD | User |
| :---: | :---: |

efdtsteq
Floating-Point Double-Precision Test Equal
efdtsteq crfD,rA,rB

$\mathrm{al} \leftarrow r \mathrm{~A}_{0: 63}$
$\mathrm{bl} \leftarrow r \mathrm{~B}_{0: 63}$
if (al = bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4 * \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is equal to $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared.
Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are generated during the execution of efdtsteq If strict IEEE-754 compliance is required, the program should use efdempeq.

Implementation note: In an implementation, the execution of efdtsteq is likely to be faster than the execution of efdcmpeq.

Floating-Point Double-Precision Test Greater Than
efdtstgt crfD,rA,rB

al $\leftarrow r A_{0: 63}$
bl $\leftarrow r B_{0: 63}$
if (al > bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4 * \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is greater than $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared.
Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are generated during the execution of efdtstgt. If strict IEEE-754 compliance is required, the program should use efdempgt.

Implementation note: In an implementation, the execution of efdtstgt is likely to be faster than the execution of efdcmpgt.

## efdtstlt

## Floating-Point Double-Precision Test Less Than

efdtstlt
crfD,rA,rB

al $\leftarrow r A_{0: 63}$
$\mathrm{bl} \leftarrow \mathrm{rB}_{0: 63}$
if (al < bl) then $c l \leftarrow 1$
else $\mathrm{cl} \leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{crD}: 4 * \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
$\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the $\mathbf{c r f D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are generated during the execution of efdtstlt. If strict IEEE-754 compliance is required, the program should use efdemplt.

Implementation note: In an implementation, the execution of efdtstlt is likely to be faster than the execution of efdemplt.

Floating-Point Absolute Value
efsabs rD,rA

$r D_{32: 63} \leftarrow 0 \mathrm{b0}| | r A_{33: 63}$
The sign bit of $\mathbf{r A}$ is cleared and the result is placed into $\mathbf{r D}$.
It is implementation dependent if invalid values for rA ( NaN , denorm, infinity) are detected and exceptions are taken.

| SPE FS | User |
| :---: | :---: |

## Floating-Point Add

efsadd
rD,rA,rB


$$
\mathrm{rD}_{32: 63} \leftarrow \mathrm{rA}_{32: 63}+{ }_{\text {sp }} \mathrm{rB}_{32: 63}
$$

The single-precision floating-point value of $\mathbf{r A}$ is added to $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$.
If an overflow condition is detected or the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are NaN or infinity, the result is an appropriately signed maximum floating-point value.

If an underflow condition is detected, the result is an appropriately signed floating-point 0 .
The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN
- FOFV if an overflow occurs
- FUNF if an underflow occurs
- FINXS, FG, FX if the result is inexact or overflow occurred and overflow exceptions are disabled


## efscfsf

SPE FS $\quad$ User

## Convert Floating-Point from Signed Fraction

efscfsf rD,rB

$r_{32: 63} \leftarrow$ CnvtI32ToFP32Sat (rB $32: 63$, SIGN, LOWER, F)
The signed fractional value in $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and placed into rD.

The following status bits are set in the SPEFSCR:

- FINXS, FG, FX if the result is inexact


## efscfsi

## Convert Floating-Point from Signed Integer

efscfsi
rD,rB


$$
r D_{32: 63} \leftarrow \text { CnvtSI32TOFP32Sat }\left(r B_{32: 63},\right. \text { SIGN, LOWER, I) }
$$

The signed integer value in $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and placed into rD.
The following status bits are set in the SPEFSCR:

- FINXS, FG, FX if the result is inexact


## Convert Floating-Point from Unsigned Fraction

efscfuf rD,rB

$r D_{32: 63} \leftarrow$ CnvtI32ToFP32Sat (rB 32:63 , UNSIGN, LOWER, F)
The unsigned fractional value in $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and placed into rD.

The following status bits are set in the SPEFSCR:

- FINXS, FG, FX if the result is inexact


## efscfui

| SPE FS | User |
| :---: | :---: |

## Convert Floating-Point from Unsigned Integer

efscfui rD,rB


The unsigned integer value in $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and placed into rD.

The following status bits are set in the SPEFSCR:

- FINXS, FG, FX if the result is inexact


## efscmpeq

| SPE FS | User |
| :---: | :---: |

## Floating-Point Compare Equal

efscmpeq
crD,rA,rB

al $\leftarrow \mathrm{rA}_{32: 63}$
bl $\leftarrow r B_{32: 63}$
if (al $=\mathrm{bl})$ then $\mathrm{cl} \leftarrow 1$
else $c l \leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4^{*} \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
The value in $\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ equals $\mathbf{r B}$, the $\mathbf{c r D}$ bit is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

If either operand contains a NaN , infinity, or a denorm and floating-point invalid exceptions are enabled in the SPEFSCR, the exception is taken. If the exception is not enabled, the comparison treats NaNs , infinities, and denorms as normalized numbers.

The following status bits are set in SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm or NaN


## Floating-Point Compare Greater Than

efscmpgt
crD,rA,rB

al $\leftarrow r A_{32: 63}$
$\mathrm{bl} \leftarrow \mathrm{rB}_{32: 63}$
if (al > bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \text { CrD }}$ 4*CrD $+3^{4}$ undefined || cl || undefined || undefined
The value in $\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is greater than $\mathbf{r} B$, the bit in the $\mathbf{c r D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

If either operand contains a NaN , infinity, or a denorm and floating-point invalid exceptions are enabled in the SPEFSCR, the exception is taken. If the exception is not enabled, the comparison treats NaNs, infinities, and denorms as normalized numbers.

The following status bits are set in SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm or NaN

| SPE FS | User |
| :---: | :---: |

Floating-Point Compare Less Than
efscmplt
crD,rA,rB

al $\leftarrow r A_{32: 63}$
bl $\leftarrow r B_{32: 63}$
if (al < bl) then $\mathrm{cl} \leftarrow 1$
else $\mathrm{cl} \leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4^{*} \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
The value in $\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the $\mathbf{c r D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

If either operand contains a NaN , infinity, or a denorm and floating-point invalid exceptions are enabled in the SPEFSCR, the exception is taken. If the exception is not enabled, the comparison treats NaNs, infinities, and denorms as normalized numbers.

The following status bits are set in SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm or NaN
SPE FS $\quad$ User


## Convert Floating-Point to Signed Fraction

efsctsf rD,rB

$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, SIGN, LOWER, ROUND, $F$ )
The single-precision floating-point value in $\mathbf{r B}$ is converted to a signed fraction using the current rounding mode. The result saturates if it cannot be represented in a 32-bit fraction. NaNs are converted to 0 .

The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r B}$ are +infinity., -infinity, denorm, or NaN , or $\mathbf{r B}$ cannot be represented in the target format
- FINXS, FG, FX if the result is inexact


## efsctsi

## Convert Floating-Point to Signed Integer

efsctsi rD,rB

$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, SIGN, LOWER, ROUND, I)
The single-precision floating-point value in $\mathbf{r B}$ is converted to a signed integer using the current rounding mode. The result saturates if it cannot be represented in a 32-bit integer. NaNs are converted to 0 .

The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN , or $\mathbf{r B}$ cannot be represented in the target format
- FINXS, FG, FX if the result is inexact


## Convert Floating-Point to Signed Integer with Round toward Zero

efsctsiz rD,rB


The single-precision floating-point value in $\mathbf{r B}$ is converted to a signed integer using the rounding mode Round towards Zero. The result saturates if it cannot be represented in a 32-bit integer. NaNs are converted to 0 .

| SPE FS | User |
| :---: | :---: |

## Convert Floating-Point to Unsigned Fraction

efsctuf rD,rB

$r D_{32: 63} \leftarrow$ CnvtFP32ToISat ( $r_{32: 63}$, UNSIGN, LOWER, ROUND, F)
The single-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned fraction using the current rounding mode. The result saturates if it cannot be represented in a 32 -bit unsigned fraction. NaNs are converted to 0 .

The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN , or $\mathbf{r B}$ cannot be represented in the target format
- FINXS, FG, FX if the result is inexact


## efsctui

| SPE FS | User |
| :---: | :---: |

## Convert Floating-Point to Unsigned Integer

efsctui rD,rB

$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB ${ }_{32: 63}$, UNSIGN, LOWER, ROUND, I)
The single-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned integer using the current rounding mode. The result saturates if it cannot be represented in a 32-bit unsigned integer. NaNs are converted to 0 .

The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN, or $\mathbf{r B}$ cannot be represented in the target format
- FINXS, FG, FX if the result is inexact


## efsctuiz

## Convert Floating-Point to Unsigned Integer with Round toward Zero

efsctuiz rD,rB

$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $\left.32: 63, ~ U N S I G N, ~ L O W E R, ~ T R U N C, ~ I\right) ~$
The single-precision floating-point value in $\mathbf{r B}$ is converted to an unsigned integer using the rounding mode Round toward Zero. The result saturates if it cannot be represented in a 32-bit unsigned integer. NaNs are converted to 0 .

The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN , or $\mathbf{r B}$ cannot be represented in the target format
- FINXS, FG, FX if the result is inexact


## efsdiv

| SPE FS | User |
| :--- | :--- |

Floating-Point Divide
efsdiv
rD,rA,rB

$\mathrm{rD}_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \div_{\mathrm{sp}} \mathrm{rB}_{32: 63}$
The single-precision floating-point value in $\mathbf{r A}$ is divided by $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$.
If an overflow is detected, or $\mathbf{r B}$ is a denorm (or 0 value), or $\mathbf{r A}$ is a NaN or infinity and $\mathbf{r B}$ is a normalized number, the result is an appropriately signed maximum floating-point value.

If an underflow is detected or rB is a NaN or infinity, the result is an appropriately signed floating-point 0 .
The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN
- FOFV if an overflow occurs
- FUNV if an underflow occurs
- FDBZS, FDBZ if a divide by zero occurs
- FINXS, FG, FX if the result is inexact or overflow occurred and overflow exceptions are disabled


## efsmul

| SPE FS | User |
| :---: | :---: |

## Floating-Point Multiply

efsmul rD,rA,rB

$r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \times_{\mathrm{sp}} \mathrm{rB}_{32: 63}$
The single-precision floating-point value in $\mathbf{r A}$ is multiplied by $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$.
If an overflow is detected the result is an appropriately signed maximum floating-point value.
If one of $\mathbf{r A}$ or $\mathbf{r B}$ is a NaN or an infinity and the other is not a denorm or zero, the result is an appropriately signed maximum floating-point value.

If an underflow is detected, or $\mathbf{r A}$ or $\mathbf{r B}$ is a denorm, the result is an appropriately signed floating-point 0 .
The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN
- FOFV if an overflow occurs
- FUNV if an underflow occurs
- FINXS, FG, FX if the result is inexact or overflow occurred and overflow exceptions are disabled


## Floating-Point Negative Absolute Value

efsnabs rD,rA

$r D_{32: 63} \leftarrow 0 \mathrm{~b} 1| | \mathrm{rA}_{33: 63}$
The sign bit of $\mathbf{r A}$ is set and the result is stored in $\mathbf{r D}$. It is implementation dependent if invalid values for $\mathbf{r A}(\mathrm{NaN}$, denorm, infinity) are detected and exceptions are taken.

## efsneg

| SPE FS | User |
| :---: | :---: |

## Floating-Point Negate

efsneg rD,rA


$$
\mathrm{rD}_{32: 63} \leftarrow \neg \mathrm{rA}_{32} \| \mathrm{rA}_{33: 63}
$$

The sign bit of $\mathbf{r A}$ is complemented and the result is stored in $\mathbf{r D}$. It is implementation dependent if invalid values for $\mathbf{r A}$ ( NaN , denorm, infinity) are detected and exceptions are taken.

| SPE FS | User |
| :---: | :---: |

Floating-Point Subtract
efssub

$$
\mathbf{r D}, \mathbf{r A , r B}
$$



$$
r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63}-\mathrm{sp} \mathrm{rB}_{32: 63}
$$

The single-precision floating-point value in $\mathbf{r B}$ is subtracted from that in $\mathbf{r A}$ and the result is stored in $\mathbf{r D}$.
If an overflow condition is detected or the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are NaN or infinity, the result is an appropriately signed maximum floating-point value.

If an underflow condition is detected, the result is an appropriately signed floating-point 0 .
The following status bits are set in the SPEFSCR:

- FINV if the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are +infinity, -infinity, denorm, or NaN
- FOFV if an overflow occurs
- FUNF if an underflow occurs
- FINXS, FG, FX if the result is inexact or overflow occurred and overflow exceptions are disabled


## efststeq

| SPE FS | User |
| :---: | :---: |

## Floating-Point Test Equal

efststeq
crD,rA,rB

al $\leftarrow \mathrm{rA}_{32: 63}$
bl $\leftarrow \mathrm{rB}_{32: 63}$
if (al $=\mathrm{bl})$ then $\mathrm{cl} \leftarrow 1$
else $c l \leftarrow 0$
$\mathrm{CR}_{4 * \mathrm{CrD}: 4^{*} \mathrm{CrD}+3} \leftarrow$ undefined ||cl || undefined || undefined
The value in $\mathbf{r A}$ is compared against $\mathbf{r B}$. If $\mathbf{r A}$ equals $\mathbf{r B}$, the bit in $\mathbf{c r D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$. The comparison treats NaNs, infinities, and denorms as normalized numbers.

No exceptions are taken during execution of efststeq. If strict IEEE-754 compliance is required, the program should use efscmpeq.
SPE FS $\quad$ User

## Floating-Point Test Greater Than

efststgt crD,rA,rB

al $\leftarrow \mathrm{rA}_{32: 63}$
bl $\leftarrow r B_{32: 63}$
if (al > bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \text { crD }: 4 * \text { crD }+3} \leftarrow$ undefined || Cl || undefined || undefined
If $\mathbf{r A}$ is greater than $\mathbf{r B}$, the bit in crD is set, otherwise it is cleared. Comparison ignores the sign of 0 $(+0=-0)$. The comparison treats NaNs, infinities, and denorms as normalized numbers.

No exceptions are taken during the execution of efststgt. If strict IEEE-754 compliance is required, the program should use efsempgt.

| SPE FS | User |
| :---: | :---: |

## Floating-Point Test Less Than

efststlt crD,rA,rB

al $\leftarrow \mathrm{rA}_{32: 63}$
$\mathrm{bl} \leftarrow \mathrm{rB}_{32: 63}$
if (al < bl) then $\mathrm{cl} \leftarrow 1$
else cl $\leftarrow 0$
$\mathrm{CR}_{4 * \text { CrD: }} \mathrm{A*CrD}^{2}+3 \leftarrow$ undefined || cl || undefined || undefined
If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the $\mathbf{c r D}$ is set, otherwise it is cleared. Comparison ignores the sign of 0 $(+0=-0)$. The comparison treats NaNs, infinities, and denorms as normalized numbers.

No exceptions are taken during the execution of efststlt. If strict IEEE-754 compliance is required, the program should use efscmplt.

## evabs

| SPE | User |
| :--- | :--- |

evabs
Vector Absolute Value
evabs rD,rA

$r D_{0: 31} \leftarrow \operatorname{ABS}\left(\mathrm{rA}_{0: 31}\right)$
$r D_{32: 63} \leftarrow \mathrm{ABS}\left(\mathrm{rA}_{32: 63}\right)$
The absolute value of each element of $\mathbf{r A}$ is placed in the corresponding elements of $\mathbf{r D}$, as shown in Figure 5-2. An absolute value of 0x8000_0000 (most negative number) returns 0x8000_0000. No overflow is detected.


Figure 5-2. Vector Absolute Value (evabs)

## evaddiw

| SPE | User |
| :--- | :--- |

## evaddiw

Vector Add Immediate Word
evaddiw rD,rB,UIMM


```
rD 0:31}\leftarrow\leftarrowr\mp@subsup{B}{0:31}{}+\mathrm{ EXTZ(UIMM) // Modulo sum
rD 32:63}\leftarrowr\mp@subsup{B}{32:63 + EXTZ (UIMM) // Modulo sum}{
```

UIMM is zero-extended and added to both the high and low elements of $\mathbf{r B}$ and the results are placed in rD, as shown in Figure 5-3. Note that the same value is added to both elements of the register. UIMM is 5 bits.


Figure 5-3. Vector Add Immediate Word (evaddiw)

## evaddsmiaaw

SPE $\quad$ User

## evaddsmiaaw

Vector Add Signed, Modulo, Integer to Accumulator Word
evaddsmiaaw rD,rA


$$
\begin{aligned}
& r D_{0: 31} \leftarrow A C C_{0: 31}+r A_{0: 31} \\
& r D_{32: 63} \leftarrow A C C_{32: 63}+r A_{32: 63} \\
& \mathrm{ACC}_{0: 63} \leftarrow r D_{0: 63}
\end{aligned}
$$

Each word element in $\mathbf{r A}$ is added to the corresponding element in the accumulator and the results are placed in rD and into the accumulator, as shown in Figure 5-4.

Other registers altered: ACC


Figure 0-1. Vector Add Signed, Modulo, Integer to Accumulator Word (evaddsmiaaw)

## evaddssiaaw

| SPE | User |
| :---: | :---: |

Vector Add Signed, Saturate, Integer to Accumulator Word
evaddssiaaw rD,rA


```
// high
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{0:31}{})+\operatorname{EXTS}(r\mp@subsup{A}{0:31}{}
ovh \leftarrowtemp 31 }\oplus\mp@subsup{\mathrm{ temp}}{32}{
rD 0:31}\leftarrowSATURATE (Ovh, temp 31, 0x80000000, Ox7ffffffff, temp 32:63) (%
// low
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{32:63}{})+\operatorname{EXTS}(r\mp@subsup{A}{32:63}{}
ovl \leftarrowtemp 31 }\oplus\mp@subsup{\mathrm{ temp}}{32}{
rD 32:63}\leftarrowSATURATE (ovl, temp 31, 0x80000000, 0x7ffffffff, temp 32:63
ACC}0:63\leftarrowr\mp@subsup{D}{0:63}{
SPEFSCR OVH 
SPEFSCR OV 
SPEFSCR 
SPEFSCR
```

Each signed integer word element in $\mathbf{r A}$ is sign-extended and added to the corresponding sign-extended element in the accumulator, saturating if overflow or underflow occurs, and the results are placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-4. Any overflow or underflow is recorded in the SPEFSCR overflow and summary overflow bits.
Other registers altered: SPEFSCR ACC


Figure 5-4. Vector Add Signed, Saturate, Integer to Accumulator Word (evaddssiaaw)

## evaddumiaaw

Vector Add Unsigned, Modulo, Integer to Accumulator Word
evaddumiaaw rD,rA


```
rD 0:31}\mp@code{\leftarrowACC0:31}+r\mp@subsup{A}{0:31}{
rD 32:63}\leftarrow~\mp@subsup{ACC}{32:63}{}+r\mp@subsup{A}{32:63}{
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

Each unsigned integer word element in $\mathbf{r A}$ is added to the corresponding element in the accumulator and the results are placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-5.

Other registers altered: ACC


Figure 5-5. Vector Add Unsigned, Modulo, Integer to Accumulator Word (evaddumiaaw)

## evaddusiaaw

Vector Add Unsigned, Saturate, Integer to Accumulator Word
evaddusiaaw rD,rA


```
// high
\mp@subsup{\operatorname{emp}}{0:63}{}\leftarrow\operatorname{EXTZ}(\mp@subsup{\textrm{ACC}}{0:31}{})+\operatorname{EXTZ}(r\mp@subsup{A}{0:31}{})
ovh }\leftarrow\mp@subsup{\mathrm{ temp }}{31}{
rD 0:31}\leftarrow\leftarrowSATURATE (ovh, temp 31, Oxffffffff, 0xfffffffff, temp 32:63
// low
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{\textrm{ACC}}{32:63}{})+\operatorname{EXTZ}(r\mp@subsup{A}{32:63}{}
ovl }\leftarrow\mp@subsup{\mathrm{ temp }}{31}{
rD 32:63}\leftarrow\mp@subsup{\mp@code{SATURATE (ovl, temp 31, 0xffffffff, Oxffffffff, temp 32:63)}}{3}{
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
SPEFSCR OvH
SPEFSCR 
SPEFSCR 
SPEFSCR 
```

Each unsigned integer word element in $\mathbf{r A}$ is zero-extended and added to the corresponding zero-extended element in the accumulator, saturating if overflow occurs, and the results are placed in rD and the accumulator, as shown in Figure 5-6. Any overflow is recorded in the SPEFSCR overflow and summary overflow bits.

Other registers altered: SPEFSCR ACC


Figure 5-6. Vector Add Unsigned, Saturate, Integer to Accumulator Word (evaddusiaaw)

## evaddw

| SPE | User |
| :--- | :--- |

Vector Add Word
evaddw
rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31}+\mathrm{rB}_{0: 31} / /$ Modulo sum
$r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63}+\mathrm{rB}_{32: 63} / /$ Modulo sum
The corresponding elements of $\mathbf{r A}$ and $\mathbf{r B}$ are added and the results are placed in $\mathbf{r D}$, as shown in Figure 5-7. The sum is a modulo sum.


Figure 5-7. Vector Add Word (evaddw)

## evand

| SPE | User |
| :--- | :--- |

Vector AND
evand
rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31} \& r \mathrm{~B}_{0: 31} / /$ Bitwise AND
$r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \& \mathrm{rB}_{32: 63} / /$ Bitwise AND

The corresponding elements of $\mathbf{r A}$ and $\mathbf{r B}$ are ANDed bitwise and the results are placed in the corresponding element of $\mathbf{r D}$, as shown in Figure 5-8.


Figure 5-8. Vector AND (evand)

## evandc

Vector AND with Complement
evandc
rD,rA,rB

$r \mathrm{D}_{0: 31} \leftarrow \mathrm{rA}_{0: 31} \&\left(\neg \mathrm{rB}_{0: 31}\right) / /$ Bitwise ANDC
$r \mathrm{D}_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \&\left(\neg \mathrm{rB}_{32: 63}\right) / /$ Bitwise ANDC
The word elements of $\mathbf{r A}$ and are ANDed bitwise with the complement of the corresponding elements of $\mathbf{r B}$. The results are placed in the corresponding element of $\mathbf{r D}$, as shown in Figure 5-9.


Figure 5-9. Vector AND with Complement (evandc)

## evcmpeq

| SPE | User |
| :--- | :--- |

Vector Compare Equal
evcmpeq crD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{B}}{0:31}{
bl }\leftarrowr\mp@subsup{B}{32:63}{
if (ah = bh) then ch }\leftarrow
else ch \leftarrow0
if (al = bl) then cl\leftarrow1
else cl \leftarrow0
CR}4*CrD:4*CrD+3 \leftarrow ch || cl || (ch | cl) || (ch & cl)
```

The most significant bit in crD is set if the high-order element of $\mathbf{r A}$ is equal to the high-order element of $\mathbf{r B}$, as shown in Figure 5-10; it is cleared otherwise. The next bit in crD is set if the low-order element of $\mathbf{r A}$ is equal to the low-order element of $\mathbf{r B}$ and cleared otherwise. The last two bits of $\mathbf{c r D}$ are set to the OR and AND of the result of the compare of the high and low elements.


Figure 5-10. Vector Compare Equal (evcmpeq)

## evcmpgts

| SPE | User |
| :---: | :---: |

## evcmpgts

Vector Compare Greater Than Signed
evempgts
crD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al }\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrow\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrow\textrm{rB}32:6
if (ah > bh) then ch \leftarrow1
else ch\leftarrow0
if (al > bl) then cl \leftarrow1
else cl }\leftarrow
CR4*CrD:4*CrD+3}\leftarrow\textrm{ch}||\textrm{cl || (ch | cl) || (ch & cl)
```

The most significant bit in crD is set if the high-order element of $\mathbf{r A}$ is greater than the high-order element of $\mathbf{r B}$, as shown in Figure 5-11; it is cleared otherwise. The next bit in crD is set if the low-order element of $\mathbf{r A}$ is greater than the low-order element of $\mathbf{r B}$ and cleared otherwise. The last two bits of $\mathbf{c r D}$ are set to the OR and AND of the result of the compare of the high and low elements.


Figure 5-11. Vector Compare Greater Than Signed (evcmpgts)

## evcmpgtu

| SPE | User |
| :---: | :---: |

Vector Compare Greater Than Unsigned
evempgtu crD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{B}}{0:31}{
bl}\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah >U bh) then ch}\leftarrow
else ch \leftarrow0
if (al >U bl) then cl\leftarrow1
else cl }\leftarrow
CR}4*\operatorname{crD}:4*\operatorname{crD}+3\leftarrow\textrm{ch}||\textrm{cl}||(\textrm{ch}|\textrm{cl})||(ch&cl
```

The most significant bit in crD is set if the high-order element of $\mathbf{r A}$ is greater than the high-order element of $\mathbf{r B}$, as shown in Figure 5-12; it is cleared otherwise. The next bit in $\mathbf{c r D}$ is set if the low-order element of $\mathbf{r A}$ is greater than the low-order element of $\mathbf{r B}$ and cleared otherwise. The last two bits of $\mathbf{c r D}$ are set to the OR and AND of the result of the compare of the high and low elements.


Figure 5-12. Vector Compare Greater Than Unsigned (evcmpgtu)

## evcmplts

| SPE | User |
| :---: | :---: |

## evcmplts

Vector Compare Less Than Signed
evemplts
crD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{B}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah < bh) then ch }\leftarrow
else ch \leftarrow0
if (al < bl) then cl \leftarrow1
else cl \leftarrow0
CR}4*\operatorname{crD}:4*\operatorname{crD}+3\leftarrow\textrm{ch}||\textrm{cl}||(\textrm{ch}|\textrm{cl})||(ch&cl
```

The most significant bit in crD is set if the high-order element of $\mathbf{r A}$ is less than the high-order element of $\mathbf{r B}$, as shown in Figure 5-13; it is cleared otherwise. The next bit in crD is set if the low-order element of $\mathbf{r A}$ is less than the low-order element of $\mathbf{r B}$ and cleared otherwise. The last two bits of $\mathbf{c r D}$ are set to the OR and AND of the result of the compare of the high and low elements.


Figure 5-13. Vector Compare Less Than Signed (evemplts)

## evcmpltu

SPE $\quad$ User

## Vector Compare Less Than Unsigned

evempltu
crD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{B}}{0:31}{
bl }\leftarrowr\mp@subsup{B}{32:63}{
if (ah <U bh) then ch\leftarrow1
else ch\leftarrow0
if (al <U bl) then cl \leftarrow1
else cl }\leftarrow
CR 4*CrD:4*crD+3
```

The most significant bit in crD is set if the high-order element of $\mathbf{r A}$ is less than the high-order element of $\mathbf{r B}$, as shown in Figure 5-14; it is cleared otherwise. The next bit in crD is set if the low-order element of $\mathbf{r A}$ is less than the low-order element of $\mathbf{r B}$ and cleared otherwise. The last two bits of $\mathbf{c r D}$ are set to the OR and AND of the result of the compare of the high and low elements.


Figure 5-14. Vector Compare Less Than Unsigned (evcmpltu)

## evcntlsw

Vector Count Leading Signed Bits Word
eventlsw rD,rA


The leading sign bits in each element of $\mathbf{r A}$ are counted, and the respective count is placed into each element of rD, as shown in Figure 5-15.
eventlzw is used for unsigned operands; eventlsw is used for signed operands.


Figure 5-15. Vector Count Leading Signed Bits Word (eventlsw)

## evcntlzw

| SPE | User |
| :--- | :--- |

## evcntlzw

Vector Count Leading Zeros Word
eventlzw rD,rA


The leading zero bits in each element of $\mathbf{r A}$ are counted, and the respective count is placed into each element of rD, as shown in Figure 5-16.


Figure 5-16. Vector Count Leading Zeros Word (eventizw)

## evdivws

| SPE | User |
| :---: | :---: |

evdivws
Vector Divide Word Signed
evdivws rD,rA,rB


```
dividendh \(\leftarrow \mathrm{rA}_{0: 31}\)
dividendl \(\leftarrow \mathrm{rA}_{32: 63}\)
divisorh \(\leftarrow r B_{0: 31}\)
divisorl \(\leftarrow r B_{32: 63}\)
\(r D_{0: 31} \leftarrow\) dividendh \(\div\) divisorh
\(r D_{32: 63} \leftarrow\) dividendl \(\div\) divisorl
ovh \(\leftarrow 0\)
ovl \(\leftarrow 0\)
if ((dividendh < 0) \& (divisorh \(=0)\) ) then
        \(r D_{0: 31} \leftarrow 0 \times 80000000\)
        ovh \(\leftarrow 1\)
else if ((dividendh >= 0) \& (divisorh \(=0\) )) then
        \(r D_{0: 31} \leftarrow 0 \times 7\) FFFFFFF
        ovh \(\leftarrow 1\)
else if ((dividendh \(=0 \times 80000000)\) \& (divisorh \(\left.=0 x F F F F \_F F F F\right)\) ) then
        \(r D_{0: 31} \leftarrow 0 \times 7\) FFFFFFF
        ovh \(\leftarrow 1\)
if ((dividendl < 0) \& (divisorl = 0)) then
        \(r D_{32}: 63 \leftarrow 0 \times 80000000\)
        ovl \(\leftarrow 1\)
else if ((dividendl >= 0) \& (divisorl = 0)) then
    \(r D_{32: 63} \leftarrow 0 \times 7\) FFFFFFF
    ovl \(\leftarrow 1\)
else if ((dividendl \(=0 \times 80000000)\) \& (divisorl \(\left.=0 x F F F F \_F F F F\right)\) ) then
        \(\mathrm{rD}_{32: 63} \leftarrow 0 \times 7\) FFFFFFF
    ovl \(\leftarrow 1\)
SPEFSCR \(_{\text {OVH }} \leftarrow\) ovh
\(\mathrm{SPEFSCR}_{\text {OV }} \leftarrow \mathrm{ovl}\)
SPEFSCR \(_{\text {SOVH }} \leftarrow\) SPEFSCR \(_{\text {SOVH }} \mid\) ovh
SPEFSCR \(_{\text {SOV }} \leftarrow\) SPEFSCR \(_{\text {SOV }} \mid\) ovl
```

The two dividends are the two elements of the rA contents. The two divisors are the two elements of the rB contents, as shown in Figure 5-17. The resulting two 32-bit quotients are placed into rD. Remainders are not supplied. The operands and quotients are interpreted as signed integers. If overflow, underflow, or divide by zero occurs, the overflow and summary overflow SPEFSCR bits are set. Note that any overflow indication is always set as a side effect of this instruction. No form is defined that disables the setting of the overflow bits. In case of overflow, a saturated value is delivered into the destination register.


Figure 5-17. Vector Divide Word Signed (evdivws)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evdivwu

| SPE | User |
| :--- | :--- |

## evdivwu

Vector Divide Word Unsigned
evdivwu rD,rA,rB


```
dividendh}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
dividendl }\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
divisorh }\leftarrow r\mp@subsup{B}{0:31}{
divisorl }\leftarrow r\mp@subsup{B}{32:63}{
rD0:31}\leftarrow < dividendh \div divisorh
rD D2:63 \leftarrow dividendl \div divisorl
ovh}\leftarrow
ovl }\leftarrow
if (divisorh = 0) then
    rD0:31 = 0xFFFFFFFF
    ovh }\leftarrow
if (divisorl = 0) then
    rD 32:63}\leftarrow0xFFFFFFF
        ovl \leftarrow }
SPEFSCR ovH }\leftarrow ov
SPEFSCR OV }\leftarrow < Ov
SPEFSCR (SOVH
SPEFSCR Sov }\leftarrow\mp@subsup{\mathrm{ SPEFSCR Sov | |ovl}}{\mathrm{ SO}}{
```

The two dividends are the two elements of the contents of $\mathbf{r A}$. The two divisors are the two elements of the contents of $\mathbf{r B}$, as shown in Figure 5-18. Two 32-bit quotients are formed as a result of the division on each of the high and low elements and the quotients are placed into rD. Remainders are not supplied. Operands and quotients are interpreted as unsigned integers. If a divide by zero occurs, the overflow and summary overflow SPEFSCR bits are set. Note that any overflow indication is always set as a side effect of this instruction. No form is defined that disables the setting of the overflow bits. In case of overflow, a saturated value is delivered into the destination register.


Figure 5-18. Vector Divide Word Unsigned (evdivwu)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## eveqv

| SPE | User |
| :--- | :--- |

eveqv
Vector Equivalent
evequ rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31} \equiv r \mathrm{~B}_{0: 31} / /$ Bitwise XNOR
$r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \equiv \mathrm{rB}_{32: 63} / /$ Bitwise XNOR
The corresponding elements of $\mathbf{r A}$ and $\mathbf{r B}$ are XNORed bitwise, and the results are placed in $\mathbf{r D}$, as shown in Figure 5-19.


Figure 5-19. Vector Equivalent (eveqv)

## evextsb

| SPE | User |
| :--- | :--- |

evextsb
Vector Extend Sign Byte
evextsb rD,rA

$r D_{0: 31} \leftarrow \operatorname{EXTS}\left(\mathrm{rA}_{24: 31}\right)$
$r D_{32: 63} \leftarrow \operatorname{EXTS}\left(\mathrm{rA}_{56: 63}\right)$
The signs of the byte in each of the elements in $\mathbf{r A}$ are extended, and the results are placed in $\mathbf{r D}$, as shown in Figure 5-20.


Figure 5-20. Vector Extend Sign Byte (evextsb)

## evextsh

| SPE | User |
| :--- | :--- |

## evextsh

Vector Extend Sign Half Word
evextsh rD,rA

$r D_{0: 31} \leftarrow \operatorname{EXTS}\left(\mathrm{rA}_{16: 31}\right)$
$r D_{32: 63} \leftarrow \operatorname{EXTS}\left(\mathrm{rA}_{48: 63}\right)$
The signs of the half words in each of the elements in $\mathbf{r A}$ are extended, and the results are placed in $\mathbf{r D}$, as shown in Figure 5-21.


Figure 5-21. Vector Extend Sign Half Word (evextsh)

## evfsabs

SPE FV
User

## evfsabs

Vector Floating-Point Single-Precision Absolute Value
evfsabs rD,rA


```
rD0:31}\mp@code{0b0 || rA 1:31
rD 32:63}\leftarrow0.00 || rA 33:63
```

The sign bit of each element in $\mathbf{r A}$ is set to 0 and the results are placed into $\mathbf{r D}$.

## Exceptions:

Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the computation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: if the contents of either element of rA are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled, an interrupt is taken and the destination register is not updated.

## evfsadd

| SPE FV | User |
| :---: | :---: |

## Vector Floating-Point Single-Precision Add

evfsadd rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31}+{ }_{\mathrm{sp}} r B_{0: 31} \\
& r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63}+_{\mathrm{sp}} r B_{32: 63}
\end{aligned}
$$

Each single-precision floating-point element of $\mathbf{r A}$ is added to the corresponding element of $\mathbf{r B}$ and the results are stored in $\mathbf{r D}$. If an element of $\mathbf{r A}$ is NaN or infinity, the corresponding result is either pmax $\left(a_{\text {sign }}=0\right)$, or $n \max \left(a_{\text {sign }}==1\right)$. Otherwise, if an element of $\mathbf{r B}$ is NaN or infinity, the corresponding result is either $\operatorname{pmax}\left(\mathrm{b}_{\text {sign }}==0\right)$, or $\operatorname{mmax}\left(\mathrm{b}_{\text {sign }}==1\right)$. Otherwise, if an overflow occurs, pmax or $n \max$ (as appropriate) is stored in the corresponding element of $\mathbf{r D}$. If an underflow occurs, +0 (for rounding modes $\mathrm{RN}, \mathrm{RZ}, \mathrm{RP}$ ) or -0 (for rounding mode RM ) is stored in the corresponding element of $\mathbf{r D}$.

## Exceptions:

If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF,FOVFH] are set appropriately, or if an underflow occurs, SPEFSCR[FUNF,FUNFH] are set appropriately. If either underflow or overflow exceptions are enabled and a corresponding status bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If either result element of this instruction is inexact, or overflows but overflow exceptions are disabled, and no other interrupt is taken, or underflows but underflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS,FINXSH] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

FG and FX (FGH and FXH) are cleared if an overflow or underflow interrupt is taken, or if an invalid operation/input error is signaled for the low (high) element (regardless of FINVE).
SPEFV User

Vector Convert Floating-Point Single-Precision from Signed Fraction
evfscfsf rD,rB

$r D_{0: 31} \leftarrow$ CnvtI32ToFP32Sat (rB $0: 31$, SIGN, UPPER, $F$ )
$r D_{32: 63} \leftarrow$ CnvtI32ToFP32Sat (rB $32: 63$, SIGN, LOWER, F)
Each signed fractional element of $\mathbf{r B}$ is converted to a single-precision floating-point value using the current rounding mode and the results are placed into the corresponding elements of $\mathbf{r D}$.

## Exceptions:

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversions are not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
SPE FV $\quad$ User

Vector Convert Floating-Point Single-Precision from Signed Integer
evfscfsi rD,rB


```
rD0:31}\leftarrow<CnvtSI32ToFP32Sat(r\mp@subsup{B}{0:31, SIGN, UPPER, I)}{
rD 32:63}\leftarrow~\mathrm{ CnvtSI32ToFP32Sat(rB 32:63, SIGN, LOWER, I)
```

Each signed integer element of $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and the results are placed into the corresponding element of $\mathbf{r D}$.

## Exceptions:

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversions are not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfscfuf

SPEFV $\quad$ User

Vector Convert Floating-Point Single-Precision from Unsigned Fraction
evfscfuf rD,rB

$r D_{0: 31} \leftarrow$ CnvtI32ToFP32Sat (rB $0: 31$, UNSIGN, UPPER, $F$ )
$r D_{32: 63} \leftarrow$ CnvtI32ToFP32Sat ( $r_{32: 63}$, UNSIGN, LOWER, F)
Each unsigned fractional element of $\mathbf{r B}$ is converted to a single-precision floating-point value using the current rounding mode and the results are placed into the corresponding elements of $\mathbf{r D}$.

Exceptions:
This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversions are not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
SPEFV $\quad$ User

Vector Convert Floating-Point Single-Precision from Unsigned Integer
evfscfui rD,rB

$r D_{0: 31} \leftarrow$ CnvtI32ToFP32Sat (rB ${ }_{031}$, UNSIGN, UPPER, I)
$r D_{32: 63} \leftarrow$ CnvtI32ToFP32Sat (rB $\mathrm{B}_{32: 63}$, UNSIGN, LOWER, I)
Each unsigned integer element of $\mathbf{r B}$ is converted to the nearest single-precision floating-point value using the current rounding mode and the results are placed into the corresponding elements of $\mathbf{r D}$.

## Exceptions:

This instruction can signal an inexact status and set SPEFSCR[FINXS] if the conversions are not exact. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfscmpeq

Vector Floating-Point Single-Precision Compare Equal
evfscmpeq crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah = bh) then ch \leftarrow 1
else ch }\leftarrow
if (al = bl) then cl \leftarrow 1
else cl }\leftarrow
CR4*crD:4*crD+3}\leftarrow ch || cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared against the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ equals $\mathbf{r B}$, the $\mathbf{c r f D}$ bit is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

## Exceptions:

If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled, an interrupt is taken, and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## evfscmpgt

SPE FV $\quad$ User

## Vector Floating-Point Single-Precision Compare Greater Than

evfscmpgt crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah > bh) then ch \leftarrow 1
else ch \leftarrow 0
if (al > bl) then cl \leftarrow 1
else cl \leftarrow 0
CR4*crD:4*crD+3}\leftarrow ch || cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared against the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ is greater than $\mathbf{r B}$, the bit in the crfD is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

Exceptions:
If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled then an interrupt is taken, and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## evfscmplt

SPE FV $\quad$ User

## Vector Floating-Point Single-Precision Compare Less Than

evfscmplt crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah < bh) then ch \leftarrow 1
else ch \leftarrow 0
if (al < bl) then cl \leftarrow 1
else cl \leftarrow 0
CR4*crD:4*CrD+3}\leftarrow~\textrm{ch}||\textrm{cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared against the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the crfD is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$.

Exceptions:
If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled then an interrupt is taken, and the condition register is not updated. Otherwise, the comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

## evfsctsf

SPE FV User

Vector Convert Floating-Point Single-Precision to Signed Fraction
evfsctsf rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB $\left.{ }_{0: 31}, ~ S I G N, ~ U P P E R, ~ R O U N D, ~ F\right)$
$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, SIGN, LOWER, ROUND, F)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to a signed fraction using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit signed fraction. NaNs are converted as though they were zero.
Exceptions:
If either element of $\mathbf{r B}$ is infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV,FINVH] are set appropriately and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfsctsi

| SPE FV | User |
| :---: | :---: |

## evfsctsi

Vector Convert Floating-Point Single-Precision to Signed Integer
evfsctsi rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB ${ }_{0: 31}$, SIGN, UPPER, ROUND, I)
$\mathrm{rD}_{32: 63} \leftarrow$ CnvtFP32ToISat ( $\mathrm{rB}_{32: 63}$, SIGN, LOWER, ROUND, I)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to a signed integer using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If the contents of either element of $\mathbf{r B}$ are infinity, denorm, or NaN , or if an overflow occurs on conversion, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

Vector Convert Floating-Point Single-Precision to Signed Integer with Round toward Zero evfsctsiz rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB $\left.{ }_{0: 31}, S I G N, ~ U P P E R, ~ T R U N C, ~ I\right)$
rD $_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, SIGN, LOWER, TRUNC, I)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to a signed integer using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

Exceptions:
If either element of $\mathbf{r B}$ is infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfsctuf

SPE FV $\quad$ User

## evfsctuf

Vector Convert Floating-Point Single-Precision to Unsigned Fraction
evfsctuf rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB $0: 31$, UNSIGN, UPPER, ROUND, $F$ )
$\mathrm{rD}_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, UNSIGN, LOWER, ROUND, F)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to an unsigned fraction using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit fraction. NaNs are converted as though they were zero.
Exceptions:
If either element of $\mathbf{r B}$ is infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.
SPEFV User

Vector Convert Floating-Point Single-Precision to Unsigned Integer
evfsctui rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB $0: 31$, UNSIGN, UPPER, ROUND, I)
$r D_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, UNSIGN, LOWER, ROUND, I)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to an unsigned integer using the current rounding mode and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

## Exceptions:

If either element of $\mathbf{r B}$ is infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfsctuiz

SPE FV
User

## evfsctuiz

Vector Convert Floating-Point Single-Precision to Unsigned Integer with Round toward Zero evfsctuiz rD,rB

$r D_{0: 31} \leftarrow$ CnvtFP32ToISat (rB $\left.0: 31, ~ U N S I G N, ~ U P P E R, ~ T R U N C, ~ I\right) ~$
$r_{32: 63} \leftarrow$ CnvtFP32ToISat (rB $32: 63$, UNSIGN, LOWER, TRUNC, I)
Each single-precision floating-point element in $\mathbf{r B}$ is converted to an unsigned integer using the rounding mode Round toward Zero and the result is saturated if it cannot be represented in a 32-bit integer. NaNs are converted as though they were zero.

Exceptions:
If either element of $\mathbf{r B}$ is infinity, denorm, or NaN, or if an overflow occurs, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken, the destination register is not updated, and no other status bits are set.

If either result element of this instruction is inexact and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result. The FGH, FXH, FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

## evfsdiv

| SPEFV | User |
| :--- | :--- |

Vector Floating-Point Single-Precision Divide
evfsdiv rD,rA,rB


```
rD}\mp@subsup{D}{0:31}{}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{}\div\mp@subsup{}{\textrm{sp}}{}\mp@subsup{\textrm{rB}}{0:31}{
rD}\mp@subsup{D}{32:63}{}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{}\div\mp@subsup{}{\textrm{sp}}{}\mp@subsup{\textrm{rB}}{32:63}{
```

Each single-precision floating-point element of $\mathbf{r A}$ is divided by the corresponding element of $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$. If an element of $\mathbf{r B}$ is a NaN or infinity, the corresponding result is a properly signed zero. Otherwise, if an element of $\mathbf{r B}$ is a zero (or a denormalized number optionally transformed to zero by the implementation), or if an element of $\mathbf{r A}$ is either NaN or infinity, the corresponding result is either $\operatorname{pmax}\left(\mathrm{a}_{\text {sign }}=\mathrm{b}_{\text {sign }}\right)$, or $\operatorname{nmax}\left(\mathrm{a}_{\text {sign }}!=\mathrm{b}_{\text {sign }}\right)$. Otherwise, if an overflow occurs, pmax or nmax (as appropriate) is stored in the corresponding element of $\mathbf{r D}$. If an underflow occurs, +0 or -0 (as appropriate) is stored in the corresponding element of $\mathbf{r D}$.

## Exceptions:

If the contents of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN , or if both $\mathbf{r A}$ and $\mathbf{r B}$ are $\pm 0$, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken and the destination register is not updated. Otherwise, if the content of $\mathbf{r B}$ is $\pm 0$ and the content of $\mathbf{r A}$ is a finite normalized non-zero number, SPEFSCR[FDBZ,FDBZH] are set appropriately. If floating-point divide-by-zero exceptions are enabled, an interrupt is then taken. Otherwise, if an overflow occurs, SPEFSCR[FOVF,FOVFH] are set appropriately, or if an underflow occurs, SPEFSCR[FUNF,FUNFH] are set appropriately. If either underflow or overflow exceptions are enabled and a corresponding bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If either result element of this instruction is inexact, or overflows but overflow exceptions are disabled, and no other interrupt is taken, or underflows but underflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

FG and FX (FGH and FXH) are cleared if an overflow or underflow interrupt is taken, or if an invalid operation/input error is signaled for the low (high) element (regardless of FINVE).

## evfsmul

SPE FV
User

## evfsmul

Vector Floating-Point Single-Precision Multiply
evfsmul rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31} \times_{\mathrm{sp}} \mathrm{rB}_{0: 31} \\
& r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63} \times_{\mathrm{sp}} \mathrm{rB}_{32: 63}
\end{aligned}
$$

Each single-precision floating-point element of $\mathbf{r A}$ is multiplied with the corresponding element of $\mathbf{r B}$ and the result is stored in $\mathbf{r D}$. If an element of $\mathbf{r A}$ or $\mathbf{r B}$ are either zero (or a denormalized number optionally transformed to zero by the implementation), the corresponding result is a properly signed zero. Otherwise, if an element of $\mathbf{r A}$ or $\mathbf{r B}$ are either NaN or infinity, the corresponding result is either $p \max \left(\mathrm{a}_{\text {sign }}==\mathrm{b}_{\text {sign }}\right)$, or $\operatorname{nmax}\left(\mathrm{a}_{\text {sign }}!=\mathrm{b}_{\text {sign }}\right)$. Otherwise, if an overflow occurs, pmax or $n \max$ (as appropriate) is stored in the corresponding element of $\mathbf{r D}$. If an underflow occurs, +0 or -0 (as appropriate) is stored in the corresponding element of $\mathbf{r D}$.

## Exceptions:

If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF,FOVFH] are set appropriately, or if an underflow occurs, SPEFSCR[FUNF,FUNFH] are set appropriately. If either underflow or overflow exceptions are enabled and a corresponding status bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If either result element of this instruction is inexact, or overflows but overflow exceptions are disabled, and no other interrupt is taken, or underflows but underflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

FG and FX (FGH and FXH) are cleared if an overflow or underflow exception is taken, or if an invalid operation/input error is signaled for the low (high) element (regardless of FINVE).

## evfsnabs

SPE FV $\quad$ User

## evfsnabs

## Vector Floating-Point Single-Precision Negative Absolute Value

evfsnabs rD,rA


```
rD 0:31}\mp@code{\leftarrow0.b1 || rA 1:31
rD 32:63}\leftarrow~0.b1 || rA A33:63
```

The sign bit of each element in $\mathbf{r A}$ is set to 1 and the results are placed into $\mathbf{r D}$.

## Exceptions:

Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the sign bit operation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: if the contents of either element of $\mathbf{r A}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled then an interrupt is taken, and the destination register is not updated.

## evfsneg

SPE FV $\quad$ User

Vector Floating-Point Single-Precision Negate
evfsneg rD,rA


$$
\begin{aligned}
& \mathrm{rD}_{0: 31} \leftarrow \neg \mathrm{rA}_{0}| | \mathrm{rA}_{1: 31} \\
& r \mathrm{D}_{32: 63} \leftarrow \neg \mathrm{rA}_{32}| | \mathrm{rA}_{33: 63}
\end{aligned}
$$

The sign bit of each element in $\mathbf{r A}$ is complemented and the results are placed into $\mathbf{r D}$.

## Exceptions:

Exception detection for embedded floating-point absolute value operations is implementation dependent. An implementation may choose to not detect exceptions and carry out the sign bit operation. If the implementation does not detect exceptions, or if exception detection is disabled, the computation can be carried out in one of two ways, as a sign bit operation ignoring the rest of the contents of the source register, or by examining the input and appropriately saturating the input prior to performing the operation.

If an implementation chooses to handle exceptions, the exception is handled as follows: if the contents of either element of rA are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If floating-point invalid input exceptions are enabled then an interrupt is taken, and the destination register is not updated.

## evfssub

SPE FV User

Vector Floating-Point Single-Precision Subtract
evfssub rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31}-{ }_{\mathrm{sp}} \mathrm{rB}_{0: 31} \\
& r D_{32: 63} \leftarrow \mathrm{rA}_{32: 63}-\mathrm{sp} \mathrm{rB}_{32: 63}
\end{aligned}
$$

Each single-precision floating-point element of $\mathbf{r B}$ is subtracted from the corresponding element of $\mathbf{r A}$ and the results are stored in $\mathbf{r D}$. If an element of $\mathbf{r A}$ is NaN or infinity, the corresponding result is either pmax $\left(a_{\text {sign }}==0\right)$, or $n \max \left(a_{\text {sign }}==1\right)$. Otherwise, if an element of $\mathbf{r B}$ is NaN or infinity, the corresponding result is either $\max \left(\mathrm{b}_{\text {sign }}==0\right)$, or $\operatorname{pmax}\left(\mathrm{b}_{\text {sign }}==1\right)$. Otherwise, if an overflow occurs, pmax or $n \max$ (as appropriate) is stored in the corresponding element of $\mathbf{r D}$. If an underflow occurs, +0 (for rounding modes $\mathrm{RN}, \mathrm{RZ}, \mathrm{RP}$ ) or -0 (for rounding mode RM ) is stored in the corresponding element of $\mathbf{r D}$.

## Exceptions:

If the contents of either element of $\mathbf{r A}$ or $\mathbf{r B}$ are infinity, denorm, or NaN, SPEFSCR[FINV,FINVH] are set appropriately, and SPEFSCR[FGH,FXH,FG,FX] are cleared appropriately. If SPEFSCR[FINVE] is set, an interrupt is taken and the destination register is not updated. Otherwise, if an overflow occurs, SPEFSCR[FOVF,FOVFH] are set appropriately, or if an underflow occurs, SPEFSCR[FUNF,FUNFH] are set appropriately. If either underflow or overflow exceptions are enabled and a corresponding status bit is set, an interrupt is taken. If any of these interrupts are taken, the destination register is not updated.

If either result element of this instruction is inexact, or overflows but overflow exceptions are disabled, and no other interrupt is taken, or underflows but underflow exceptions are disabled, and no other interrupt is taken, SPEFSCR[FINXS] is set. If the floating-point inexact exception is enabled, an interrupt is taken using the floating-point round interrupt vector. In this case, the destination register is updated with the truncated result(s). The FG and FX bits are properly updated to allow rounding to be performed in the interrupt handler.

FG and FX (FGH and FXH) are cleared if an overflow or underflow interrupt is taken, or if an invalid operation/input error is signaled for the low (high) element (regardless of FINVE).

## evfststeq

Vector Floating-Point Single-Precision Test Equal
evfststeq crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah = bh) then ch }\leftarrow 
else ch }\leftarrow
if (al = bl) then cl }\leftarrow
else cl }\leftarrow
CR4*crD:4*crD+3}\leftarrow ch || cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared against the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ equals $\mathbf{r B}$, the bit in $\mathbf{c r f D}$ is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are taken during the execution of evfststeq. If strict IEEE-754 compliance is required, the program should use evfscmpeq.

Implementation note: In an implementation, the execution of evfststeq is likely to be faster than the execution of evfscmpeq.

## evfststgt

SPE FV $\quad$ User
evfststgt
Vector Floating-Point Single-Precision Test Greater Than
evfststgt crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah > bh) then ch \leftarrow 1
else ch }\leftarrow
if (al > bl) then cl \leftarrow 1
else cl }\leftarrow
CR4*crD:4*crD+3}\leftarrow ch || cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared against the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ is greater than $\mathbf{r B}$, the bit in crfD is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are taken during the execution of evfststgt. If strict IEEE-754 compliance is required, the program should use evfscmpgt.

Implementation note: In an implementation, the execution of evfststgt is likely to be faster than the execution of evfscmpgt.

## evfststlt

| SPEFV | User |
| :--- | :--- |

Vector Floating-Point Single-Precision Test Less Than
evfststlt crfD,rA,rB


```
ah}\leftarrow\mp@subsup{\textrm{rA}}{0:31}{
al}\leftarrow\mp@subsup{\textrm{rA}}{32:63}{
bh}\leftarrowr\mp@subsup{\textrm{rB}}{0:31}{
bl }\leftarrowr\mp@subsup{\textrm{B}}{32:63}{
if (ah < bh) then ch \leftarrow 1
else ch \leftarrow 0
if (al < bl) then cl \leftarrow 1
else cl \leftarrow 0
CR4*crD:4*CrD+3}\leftarrow ch || cl || (ch | cl) || (ch & cl)
```

Each element of $\mathbf{r A}$ is compared with the corresponding element of $\mathbf{r B}$. If $\mathbf{r A}$ is less than $\mathbf{r B}$, the bit in the crfD is set, otherwise it is cleared. Comparison ignores the sign of $0(+0=-0)$. The comparison proceeds after treating NaNs, infinities, and denorms as normalized numbers, using their values of ' $e$ ' and ' $f$ ' directly.

No exceptions are taken during the execution of evfststlt. If strict IEEE-754 compliance is required, the program should use evfscmplt.
Implementation note: In an implementation, the execution of evfststlt is likely to be faster than the execution of evfscmplt.

## evidd

| SPE, SPE FV, SPE FD | User |
| :---: | :---: |

Vector Load Double Word into Double Word
evldd rD,d(rA)


The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-22 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | C | d | e | f | g | h |

GPR in big endian | a | b | c | d | e | f | g | h |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

GPR in little endian | $h$ | g | f | e | d | c | b | a |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 5-22. evIdd Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## eviddx

| SPE, SPE FV, SPE FD | User |
| :---: | :---: |

Vector Load Double Word into Double Word Indexed
evlddx
rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA}\leftarrow\textrm{b}+(r\textrm{B}
rD \leftarrowMEM(EA, 8)
```

The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-23 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | c | d | e | f | g | h |
| GPR in big endian | a | b | c | d | e | f | g | h |
| GPR in little endian | h | g | $f$ | e | d | c | b | a |

Figure 5-23. evIddx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evidh

## Vector Load Double into Four Half Words

evldh
rD,d(rA)


The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-24 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | c | d | e | $f$ | g | h |

GPR in big endian | a | b | c | d | e | f | g | h |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

GPR in little endian | b | a | d | c | f | e | h | g |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

Figure 5-24. evldh Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evidhx

| SPE | User |
| :---: | :---: |

Vector Load Double into Four Half Words Indexed
evldhx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD 0:15 }\leftarrow\operatorname{MEM(EA, 2)
rD 16:31}\leftarrow MEM(EA+2,2
rD 32:47}\leftarrow\operatorname{MEM (EA+4,2)
rD 48:63}\leftarrow\operatorname{MEM(EA+6,2)
```

The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-25 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | C | d | e | f | g | h |

GPR in big endian | a | b | c | d | e | f | g | h |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

GPR in little endian | $b$ | $a$ | $d$ | $c$ | $f$ | $e$ | $h$ | $g$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 5-25. evidhx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evldw

| SPE | User |
| :--- | :--- |

evldw
Vector Load Double into Two Words
evldw rD,d(rA)


The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-26 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-26. evldw Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evldwx

| SPE | User |
| :---: | :---: |

Vector Load Double into Two Words Indexed
evldwx
rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA}\leftarrow\textrm{b}+(rB
rD0:31}\leftarrow\leftarrowMEM(EA, 4
rD 32:63}\leftarrow~MEM(EA+4, 4
```

The double word addressed by EA is loaded from memory and placed in rD.
Figure 5-27 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |  |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Memory | a | b | c | d | e | f | g | h |
|  |  |  |  |  |  |  |  |  |  |



Figure 5-27. evldwx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evlhhesplat

Vector Load Half Word into Half Words Even and Splat
evlhhesplat rD,d(rA)

The half word addressed by EA is loaded from memory and placed in the even half words of each element of $\mathbf{r D}$.

Figure 5-28 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.

| Byte address | 0 |  |
| ---: | ---: | ---: |
| Memory | a | b |
|  |  |  |



Figure 5-28. evlhhesplat Results in Big- and Little-Endian Modes
Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evIhhesplatx

Vector Load Half Word into Half Words Even and Splat Indexed
evlhhesplatx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow (rA
EA \leftarrow b + (rB)
rD 0:15}\leftarrow\leftarrow\operatorname{MEM (EA, 2)
rD}\mp@subsup{D}{16:31}{\leftarrow}\leftarrow0\times000
rD 32:47}\leftarrow\operatorname{MEM(EA,2)
rD 48:63}\leftarrow 0x0000
```

The half word addressed by EA is loaded from memory and placed in the even half words of each element of $\mathbf{r D}$.

Figure 5-29 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-29. evlhhesplatx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evlhhossplat

| SPE | User |
| :--- | :--- |

evIhhossplat
Vector Load Half Word into Half Word Odd Signed and Splat
evlhhossplat rD,d(rA)


The half word addressed by EA is loaded from memory and placed in the odd half words sign extended in each element of rD.

Figure 5-30 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.


Figure 5-30. evlhhossplat Results in Big- and Little-Endian Modes
In big-endian memory, the msb of a is sign extended. In little-endian memory, the msb of b is sign extended.

Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evlhhossplatx

SPE $\quad$ User

Vector Load Half Word into Half Word Odd Signed and Splat Indexed
evlhhossplatx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow (rA
EA \leftarrow b + (rB)
rD 0:31}\leftarrow\leftarrow\operatorname{EXTS}(MEM(EA,2)
rD 32:63}\leftarrow\operatorname{EXTS}(MEM(EA,2)
```

The half word addressed by EA is loaded from memory and placed in the odd half words sign extended in each element of rD.

Figure 5-31 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-31. evIhhossplatx Results in Big- and Little-Endian Modes
In big-endian memory, the msb of a is sign extended. In little-endian memory, the msb of b is sign extended.
Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evlhhousplat

SPE $\quad$ User

## Vector Load Half Word into Half Word Odd Unsigned and Splat

evlhhousplat rD,d(rA)


The half word addressed by EA is loaded from memory and placed in the odd half words zero extended in each element of $\mathbf{r D}$.

Figure 5-32 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.


Figure 5-32. evIhhousplat Results in Big- and Little-Endian Modes
Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evIhhousplatx

SPE $\quad$ User

## evlhhousplatx

Vector Load Half Word into Half Word Odd Unsigned and Splat Indexed
evlhhousplatx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD 0:15}\leftarrow~0x000
rD 16:31}\leftarrow\operatorname{MEM(EA,2)
rD 32:47}\leftarrow < 0x0000
rD 48:63}\leftarrow\operatorname{MEM (EA, 2)
```

The half word addressed by EA is loaded from memory and placed in the odd half words zero extended in each element of $\mathbf{r D}$.

Figure 5-33 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-33. evlhhousplatx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not half-word aligned, an alignment exception occurs.

## evlwhe

| SPE | User |
| :--- | :--- |

Vector Load Word into Two Half Words Even
evlwhe rD,d(rA)


The word addressed by EA is loaded from memory and placed in the even half words in each element of rD.

Figure 5-34 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | c | d |



Figure 5-34. evlwhe Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhex

Vector Load Word into Two Half Words Even Indexed
evlwhex rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD 0:15}\leftarrow\leftarrow\operatorname{MEM (EA,2)
rD}\mp@subsup{D}{16:31}{\leftarrow}\leftarrow0\times000
rD 32:47}\leftarrow\operatorname{MEM(EA+2,2)
rD}48:63\leftarrow0x000
```

The word addressed by EA is loaded from memory and placed in the even half words in each element of rD.
Figure 5-35 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.


Figure 5-35. evIwhex Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhos

| SPE | User |
| :--- | :--- |

## evlwhos

Vector Load Word into Two Half Words Odd Signed (with sign extension)
evlwhos rD,d(rA)

The word addressed by EA is loaded from memory and placed in the odd half words sign extended in each element of $\mathbf{r D}$.

Figure 5-36 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.


Figure 5-36. evlwhos Results in Big- and Little-Endian Modes
In big-endian memory, the most significant bits of a and c are sign extended. In little-endian memory, the most significant bits of $b$ and $d$ are sign extended.

Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhosx

| SPE | User |
| :--- | :--- |

evlwhosx
Vector Load Word into Two Half Words Odd Signed Indexed (with sign extension)
evlwhosx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD 0:31}\leftarrow\leftarrow\operatorname{EXTS}(MEM(EA,2)
rD 32:63}\leftarrow\operatorname{EXTS}(\operatorname{MEM}(\textrm{EA}+2,2)
```

The word addressed by EA is loaded from memory and placed in the odd half words sign extended in each element of $\mathbf{r D}$.

Figure 5-37 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-37. evlwhosx Results in Big- and Little-Endian Modes
In big-endian memory, the most significant bits of a and c are sign extended. In little-endian memory, the most significant bits of $b$ and $d$ are sign extended.
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhou

| SPE | User |
| :--- | :--- |

Vector Load Word into Two Half Words Odd Unsigned (zero-extended)
evlwhou rD,d(rA)


The word addressed by EA is loaded from memory and placed in the odd half words zero extended in each element of $\mathbf{r D}$.

Figure 5-38 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-38. evlwhou Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhoux

| SPE | User |
| :---: | :---: |

Vector Load Word into Two Half Words Odd Unsigned Indexed (zero-extended) evlwhoux rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD}0:15 \leftarrow 0x0000
rD 16:31}\leftarrow < MEM(EA,2
rD 32:47}\leftarrow 0x0000
rD 48:63}\leftarrow\operatorname{MEM(EA+2,2)
```

The word addressed by EA is loaded from memory and placed in the odd half words zero extended in each element of $\mathbf{r D}$.

Figure 5-39 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-39. evlwhoux Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhsplat

Vector Load Word into Two Half Words and Splat
evlwhsplat rD,d(rA)


The word addressed by EA is loaded from memory and placed in both the even and odd half words in each element of $\mathbf{r D}$.

Figure 5-40 shows how bytes are loaded into $\mathbf{r D}$ as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 |
| ---: | :---: | :---: | :---: | :---: |
| Memory | a | b | c | d |
|  |  |  |  |  |

GPR in big endian | a | b | a | b | c | d | c | d |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

GPR in little endian | b | a | b | a | d | c | d | c |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 5-40. evlwhsplat Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwhsplatx

Vector Load Word into Two Half Words and Splat Indexed
evlwhsplatx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
rD 0:15
rD 16:31}\leftarrow < MEM(EA,2)
rD 32:47}\leftarrowMEM(EA+2,2
rD 48:63}\leftarrow\operatorname{MEM(EA+2,2)
```

The word addressed by EA is loaded from memory and placed in both the even and odd half words in each element of $\mathbf{r D}$.

Figure 5-41 shows how bytes are loaded into rD as determined by the endian mode.

| Byte address | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| Memory | a | b | c | d |

GPR in big endian | a | b | a | b | c | d | c | d |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

GPR in little endian | $b$ | a | b | a | d | c | d | c |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |

Figure 5-41. evlwhsplatx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwwsplat

| SPE | User |
| :--- | :--- |

Vector Load Word into Word and Splat
evlwwsplat rD,d(rA)


The word addressed by EA is loaded from memory and placed in both elements of rD.
Figure 5-42 shows how bytes are loaded into rD as determined by the endian mode.


Figure 5-42. evlwwsplat Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evlwwsplatx

| SPE | User |
| :---: | :---: |

## Vector Load Word into Word and Splat Indexed

evlwwsplatx rD,rA,rB


```
if (rA = 0) then b }\leftarrow
```

else $\mathrm{b} \leftarrow(r A)$
$\mathrm{EA} \leftarrow \mathrm{b}+(r \mathrm{~B})$
$r D_{0: 31} \leftarrow \operatorname{MEM}(E A, 4)$
$r D_{32: 63} \leftarrow \operatorname{MEM}(E A, 4)$

The word addressed by EA is loaded from memory and placed in both elements of $\mathbf{r D}$.
Figure 5-43 shows how bytes are loaded into rD as determined by the endian mode.


GPR in big endian | a | b | c | d | a | b | c | d |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |



Figure 5-43. evlwwsplatx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evmergehi

| SPE, SPE FV, SPE FD | User |
| :---: | :---: |

evmergehi
Vector Merge High
evmergehi rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31}$
$r D_{32: 63} \leftarrow r B_{0: 31}$
The high-order elements of $\mathbf{r A}$ and $\mathbf{r B}$ are merged and placed into $\mathbf{r D}$, as shown in Figure 5-44.


Figure 5-44. High Order Element Merging (evmergehi)
Note: A vector splat high can be performed by specifying the same register in $\mathbf{r A}$ and $\mathbf{r B}$.

## evmergehilo

SPE, SPE FV, SPE FD $\quad$ User

Vector Merge High/Low
evmergehilo rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{0: 31}$
$r D_{32: 63} \leftarrow r B_{32: 63}$
The high-order element of $\mathbf{r A}$ and the low-order element of $\mathbf{r B}$ are merged and placed into $\mathbf{r D}$, as shown in Figure 5-45.


Figure 5-45. High Order Element Merging (evmergehilo)
Application note: With appropriate specification of $\mathbf{r A}$ and $\mathbf{r B}$, evmergehi, evmergelo, evmergehilo, and evmergelohi provide a full 32-bit permute of two source operands.

## evmergelo

Vector Merge Low
evmergelo
rD,rA,rB

$r D_{0: 31} \leftarrow \mathrm{rA}_{32: 63}$
$r D_{32: 63} \leftarrow r B_{32: 63}$
The low-order elements of $\mathbf{r A}$ and $\mathbf{r B}$ are merged and placed in $\mathbf{r D}$, as shown in Figure 5-46.


Figure 5-46. Low Order Element Merging (evmergelo)
Note: A vector splat low can be performed by specifying the same register in $\mathbf{r} A$ and $\mathbf{r B}$.

## evmergelohi

| SPE | User |
| :--- | :--- |

evmergelohi
Vector Merge Low/High
evmergelohi rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow \mathrm{rA}_{32: 63} \\
& r D_{32: 63} \leftarrow r B_{0: 31}
\end{aligned}
$$

The low-order element of $\mathbf{r A}$ and the high-order element of $\mathbf{r B}$ are merged and placed into $\mathbf{r D}$, as shown in Figure 5-47.


Figure 5-47. Low Order Element Merging (evmergelohi)
Note: A vector swap can be performed by specifying the same register in $\mathbf{r A}$ and $\mathbf{r B}$.

## evmhegsmfaa

| SPE | User |
| :--- | :--- |

evmhegsmfaa
Vector Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate
evmhegsmfaa rD,rA,rB


The corresponding low even-numbered, half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The product is added to the contents of the 64-bit accumulator and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-48.

Note: This is a modulo sum. There is no overflow check and no saturation is performed. Any overflow of the 64-bit sum is not recorded into the SPEFSCR.


Figure 5-48. evmhegsmfaa (Even Form)

## evmhegsmfan

| SPE | User |
| :---: | :---: |

## evmhegsmfan

Vector Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate Negative
evmhegsmfan rD,rA,rB



```
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ temp 0:31)}}{0}{}
rD 0:63}\leftarrow ACCC0:63 - temp 0:63
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

The corresponding low even-numbered, half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The product is subtracted from the contents of the 64-bit accumulator and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-49.

Note: This is a modulo difference. There is no overflow check and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-49. evmhegsmfan (Even Form)

## evmhegsmiaa

| SPE | User |
| :---: | :---: |

evmhegsmiaa
Vector Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate
evmhegsmiaa rD,rA,rB


The corresponding low even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended and added to the contents of the 64-bit accumulator, and the resulting sum is placed into rD and into the accumulator, as shown in Figure 5-50.

Note: This is a modulo sum. There is no overflow check and no saturation is performed. Any overflow of the 64-bit sum is not recorded into the SPEFSCR.


Figure 5-50. evmhegsmiaa (Even Form)

## evmhegsmian

| SPE | User |
| :---: | :---: |

evmhegsmian
Vector Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate Negative
evmhegsmian rD,rA,rB


The corresponding low even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended and subtracted from the contents of the 64-bit accumulator, and the result is placed into $\mathbf{r D}$ and into the accumulatorFigure 5-51.

Note: This is a modulo difference. There is no check for overflow and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-51. evmhegsmian (Even Form)

## evmhegumiaa

| SPE | User |
| :--- | :--- |

evmhegumiaa
Vector Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate
evmhegumiaa rD,rA,rB


The corresponding low even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is zero-extended and added to the contents of the 64-bit accumulator. The resulting sum is placed into rD and into the accumulator, as shown in Figure 5-52.

Note: This is a modulo sum. There is no overflow check and no saturation is performed. Any overflow of the 64-bit sum is not recorded into the SPEFSCR.


Figure 5-52. evmhegumiaa (Even Form)

## evmhegumian

Vector Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate Negative
evmhegumian rD,rA,rB

| 0 |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  |  |  |  |  |  |  | 31 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 1 | 0 | 0 | 0 |  |



```
temp 0:63}\leftarrow EXTZ (temp 0:31) ()
rD 0:63}\leftarrow ACC(0:63 - temp 0:63
// update accumulator
ACC}0:63~r\mp@subsup{D}{0:63}{
```

The corresponding low even-numbered unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is zero-extended and subtracted from the contents of the 64-bit accumulator. The result is placed into $\mathbf{r D}$ and into the accumulatorFigure 5-53.

Note: This is a modulo difference. There is no check for overflow and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-53. evmhegumian (Even Form)

## evmhesmf

Vector Multiply Half Words, Even, Signed, Modulo, Fractional (to Accumulator)

| evmhesmf | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmhesmfa | rD,rA,rB | $(A=1)$ |


| 0 |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  | 25 | 26 | 27 |  |  |  | 31 | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 0 | 0 | 0 | A | 0 | 1 | 0 | 1 | 1 | 1 |

```
// high
rD 0:31
// low
rD 32:63}\leftarrow~(r\mp@subsup{A}{32:47}{}\mp@subsup{\times}{\mathrm{ sf }}{}r\mp@subsup{r}{32:47}{}
// update accumulator
if A = 1 then ACC 0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied then placed into the corresponding words of rDFigure 5-54.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-54. Even Multiply of Two Signed Modulo Fractional Elements (to Accumulator) (evmhesmf)

## evmhesmfaaw

| SPE | User |
| :--- | :--- |

## evmhesmfaaw

Vector Multiply Half Words, Even, Signed, Modulo, Fractional and Accumulate into Words evmhesmfaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrow(r\mp@subsup{A}{0:15}{}}\mp@subsup{\times}{\mathrm{ sf }}{}\mp@subsup{rB}{0:15}{}
rD 0:31}\mp@code{\leftarrow ACC0:31 + temp 0:31
// low
temp 0:31}\leftarrow\leftarrow(r\mp@subsup{A}{32:47 }{* }\mp@subsup{X}{\mathrm{ Sf }}{}r\mp@subsup{\textrm{B}}{32:47}{}
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{t}{}{\prime2mp}0:3
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32 bits of each intermediate product are added to the contents of the accumulator words to form intermediate sums, which are placed into the corresponding rD words and into the accumulator, as shown in Figure 5-55.

Other registers altered: ACC


Figure 5-55. Even Form of Vector Half-Word Multiply (evmhesmfaaw)

## evmhesmfanw

Vector Multiply Half Words, Even, Signed, Modulo, Fractional and Accumulate Negative into Words
evmhesmfanw rD,rA,rB

| 0 |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  |  |  |  |  |  |  | 31 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 |

```
// high
temp}0:31 \leftarrowr\mp@subsup{r}{0:15}{}\mp@subsup{\times}{\mathrm{ sf }}{}\mp@subsup{rBB}{0:15}{
rD0:31}\leftarrow~\mp@subsup{ACC}{0:31 - temp0:31}{0
// low
tempo:31}\leftarrow\mp@subsup{rAA}{32:47 }{\mp@subsup{X}{\mathrm{ sf }}{}}\mp@subsup{r~B}{32:47}{
rD 32:63}\leftarrow\mp@subsup{ACC}{32:63 - temp 0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32-bit intermediate products are subtracted from the contents of the accumulator words to form intermediate differences, which are placed into the corresponding rD words and into the accumulator, as shown in Figure 5-56.

Other registers altered: ACC


Figure 5-56. Even Form of Vector Half-Word Multiply (evmhesmfanw)

## evmhesmi

Vector Multiply Half Words, Even, Signed, Modulo, Integer (to Accumulator)

| evmhesmi | rD,rA,rB | $(\mathbf{A}=\mathbf{0})$ |
| :--- | :--- | :--- |
| evmhesmia | rD,rA,rB | $(A=1)$ |



```
// high
rD 0:31}\mp@code{rA\mp@subsup{A}{0:15 }{* }
// low
rD 32:63}\leftarrow~r\mp@subsup{A}{32:47}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{r}{32:47}{
// update accumulator
if A = 1, then ACC 0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The two 32-bit products are placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-57.
If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-57. Even Form for Vector Multiply (to Accumulator) (evmhesmi)

## evmhesmiaaw

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate into Words
evmhesmiaaw rD,rA,rB


```
// high
\mp@subsup{emp}{0:31}{}\leftarrowr\mp@subsup{A}{0:15}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{B}{0:15}{}
rD 0:31}\leftarrow\leftarrowACC0:31 + \mp@subsup{temp}{0:31}{0:3
// low
tempo:31}\leftarrow~\mp@subsup{r}{32:47 }{~}\mp@subsup{X}{\mathrm{ Si }}{}r\mp@subsup{B}{32:47}{
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{\mathrm{ temp}}{0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate 32-bit product is added to the contents of the accumulator words to form intermediate sums, which are placed into the corresponding rD words and into the accumulator, as shown in Figure 5-58.
Other registers altered: ACC


Figure 5-58. Even Form of Vector Half-Word Multiply (evmhesmiaaw)

## evmhesmianw

| SPE | User |
| :---: | :---: |

evmhesmianw
Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate Negative into Words
evmhesmianw rD,rA,rB


```
// high
temp00:31 \leftarrowrA 0:15 }\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{rBB}{0:15}{0
rD 0:31}\leftarrow~ACC0:31 - temp00:31
// low
temp1 0:31}\leftarrow~r\mp@subsup{A}{32:47}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{\mp@code{B}}{32:47}{
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63 - temp10:31}{
// update accumulator
ACC 0:63}\leftarrow r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate 32-bit product is subtracted from the contents of the accumulator words to form intermediate differences, which are placed into the corresponding $\mathbf{r D}$ words and into the accumulator, as shown in Figure 5-59.

Other registers altered: ACC


Figure 5-59. Even Form of Vector Half-Word Multiply (evmhesmianw)

## evmhessf

Vector Multiply Half Words, Even, Signed, Saturate, Fractional (to Accumulator)

| evmhessf | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmhessfa | rD,rA,rB | $(A=1)$ |



```
// high
tempo:31}\mp@code{rA 0:15 积f r\mp@subsup{B}{0:15}{}
if (rA 0:15 = 0x8000) & (rB 0:15 = 0x8000) then
    rD 0:31 \leftarrow0x7FFF_FFFF //saturate
    movh \leftarrow1
else
    rD0:31}\leftarrow\mp@subsup{\mp@code{temp}0:31}{0}{0
    movh }\leftarrow
// low
tempo:31
if (rA 32:47 = 0x8000) & (rB 32:47 = 0x8000) then
    rD 32:63}\leftarrow0\times7FFF_FFFF //saturate
    movl \leftarrow1
else
    rD 32:63}\mp@code{\leftarrowtemp 0:31
    movl \leftarrow0
// update accumulator
if A = 1 then ACC0:63 \leftarrowrD0:63
// update SPEFSCR
SPEFSCR ovH }\leftarrow\mathrm{ movh
SPEFSCROv }\leftarrow\mathrm{ movl
SPEFSCR SOVH}<<\mp@subsup{\mathrm{ SPEFSCR }}{\mathrm{ SOVH | | movh}}{
SPEFSCR
```

The corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32 bits of each product are placed into the corresponding words of rD, as shown in Figure 5-60. If both inputs are -1.0 , the result saturates to the largest positive signed fraction and the overflow and summary overflow bits are recorded in the SPEFSCR.
If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: $\quad$ SPEFSCR, ACC (If A = 1)


Figure 5-60. Even Multiply of Two Signed Saturate Fractional Elements (to Accumulator) (evmhessf)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhessfaaw

Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate into Words
evmhessfaaw rD,rA,rB


The corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32 -bit product. If both inputs are -1.0 , the result saturates to 0x7FFF_FFFF. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow or underflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-61.

If there is an overflow or underflow from either the multiply or the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-61. Even Form of Vector Half-Word Multiply (evmhessfaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate Negative into Words

```
evmhessfanw rD,rA,rB
```



```
// high
tempo:31}\mp@code{rA 0:15 积f r\mp@subsup{B}{0:15}{}
if (rA 0:15 = 0x8000) & (rB0:15 = 0x8000) then
    tempo:31}\leftarrow0x7FFF_FFFF //saturate
    movh }\leftarrow
else
    movh }\leftarrow
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC 0:31 }}{0}{\prime})-\operatorname{EXTS}(\mp@subsup{\mathrm{ temp}}{0:31}{}
ovh }\leftarrow(\mp@subsup{\mathrm{ temp }}{31}{}\oplus\mp@subsup{\mathrm{ temp}}{32}{}
rD 0:31 }\leftarrow\mathrm{ SATURATE (ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// low
```



```
if (rA 32:47 = 0x8000) & (rB 32:47 = 0x8000) then
    temp 0:31}\leftarrow0x7FFF_FFFF //saturate
    movl \leftarrow1
else
    movl }\leftarrow
temp 0:63 \leftarrowEXTS (ACC 32:63) - EXTS (temp 0:31)
ovl }\leftarrow(\mp@subsup{\mathrm{ temp}}{31}{}\oplus\mp@subsup{\mathrm{ temp }}{32}{\prime}
rD 32:63}\leftarrow\mathrm{ SATURATE (ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCROVH}\leftarrow\mp@code{movh
SPEFSCR 
SPEFSCR SOVH }\leftarrow\mp@subsup{\mathrm{ SPEFSCR }}{\mathrm{ SOVH }}{|}| ovh | movh
SPEFSCR SOv }\leftarrow\mp@subsup{\mathrm{ SPEFSCRROV | ovl| movl}}{\mathrm{ SOV }}{
```

The corresponding even-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32 -bit product. If both inputs are -1.0 , the result saturates to $0 x 7$ FFF_FFFF. Each 32-bit product is then subtracted from the corresponding word in the accumulator, saturating if overflow or underflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-62.

If there is an overflow or underflow from either the multiply or the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-62. Even Form of Vector Half-Word Multiply (evmhessfanw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhessiaaw

SPE User

Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate into Words
evmhessiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{0:15}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{r}{0:15}{}
```



```
ovh }\leftarrow(\mp@subsup{\mathrm{ temp 31 }}{}{\prime}\mp@subsup{\mathrm{ temp 32}}{2}{\prime}
rD 0:31}\leftarrow~SATURATE (OVh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63
// low
```



```
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{32:63)}{})+\operatorname{EXTS}(temp 0:31
ovl \leftarrow }\leftarrow(\mp@subsup{\mathrm{ temp}}{31}{}\oplus\mp@subsup{\mathrm{ temp }}{32}{}
rD 32:63}\leftarrowSATURATE (OVl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrow~\mp@subsup{rD}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow\mathrm{ ovh
SPEFSCR OV }\leftarrow\mathrm{ ovl
SPEFSCR 
SPEFSCR
```

The corresponding even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-63.

If there is an overflow or underflow from either the multiply or the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-63. Even Form of Vector Half-Word Multiply (evmhessiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhessianw

| SPE | User |
| :--- | :--- |

## evmhessianw

Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate Negative into Words
evmhessianw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{0:15}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{r}{0:15}{}
\mp@subsup{temp}{0:63}{\leftarrowEXTS (ACC 0:31) - EXTS (temp 0:31)}
ovh }\leftarrow(\mp@subsup{\mathrm{ temp 31 }}{}{\prime}\mp@subsup{\mathrm{ temp 32}}{2}{\prime}
rD 0:31 \leftarrowSATURATE (Ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// low
```



```
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\operatorname{ACC}}{32:63)}{})-\operatorname{EXTS}(\mp@subsup{t}{\mathrm{ memp 0:31}}{0}
ovl \leftarrow }\leftarrow(\mp@subsup{\mathrm{ temp}}{31}{}\oplus\mp@subsup{\mathrm{ temp 32 )}}{3}{\prime
rD 32:63}\leftarrowSATURATE (Ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow\mathrm{ Ovh
SPEFSCR OV }\leftarrow\mathrm{ ovl
SPEFSCR 
SPEFSCR
```

The corresponding even-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then subtracted from the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-64.

If there is an overflow or underflow from either the multiply or the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-64. Even Form of Vector Half-Word Multiply (evmhessianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmheumi

| SPE | User |
| :---: | :---: |

## evmheumi

Vector Multiply Half Words, Even, Unsigned, Modulo, Integer (to Accumulator)
evmheumi
rD,rA,rB
( $\mathrm{A}=0$ )
evmheumia
rD,rA,rB
( $\mathrm{A}=1$ )


```
// high
```



```
// low
```



```
// update accumulator
if A = 1 then ACC 0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The two 32-bit products are placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-65.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.


Figure 5-65. Vector Multiply Half Words, Even, Unsigned, Modulo, Integer (to Accumulator) (evmheumi)

## evmheumiaaw

| SPE | User |
| :---: | :---: |

## evmheumiaaw

Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate into Words
evmheumiaaw rD,rA,rB


```
// high
tempo:31}\leftarrow\leftarrowr\mp@subsup{A}{0:15}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rB}{0:15}{0
rD0:31}\leftarrow\leftarrowACC0:31 + temp 0:31
// low
tempo:31}\leftarrowr\mp@subsup{A}{32:47 }{~}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{32:47}{
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{\mathrm{ temp}}{0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate product is added to the contents of the corresponding accumulator words and the sums are placed into the corresponding rD and accumulator words, as shown in Figure 5-66.
Other registers altered: ACC


Figure 5-66. Even Form of Vector Half-Word Multiply (evmheumiaaw)

## evmheumianw

| SPE | User |
| :---: | :---: |

Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate Negative into Words
evmheumianw rD,rA,rB


```
// high
\mp@subsup{temp 0:31}{}{~r\mp@subsup{A}{0:15}{}}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rBB}{0:15}{}
rD0:31}\leftarrow\leftarrowACC0:31 - temp0:31
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63 - temp 0:31}{
// update accumulator
ACC
```

For each word element in the accumulator, the corresponding even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate product is subtracted from the contents of the corresponding accumulator words. The differences are placed into the corresponding $\mathbf{r D}$ and accumulator words, as shown in Figure 5-67.

Other registers altered: ACC


Figure 5-67. Even Form of Vector Half-Word Multiply (evmheumianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmheusiaaw

| SPE | User |
| :---: | :---: |

Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate into Words
evmheusiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{0:15}{}}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{0:15}{
\mp@subsup{temp}{0:63}{\leftarrowE EXTZ (ACC 0:31) + EXTZ (temp 0:31)}
ovh }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD 0:31}\leftarrow \leftarrowATURATE (Ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp 32:63
//low
```



```
temp 0:63 \leftarrow EXTZ (ACC 32:63) + EXTZ(temp 0:31)
ovl \leftarrow temp 31
rD 32:63 \leftarrow SATURATE(ovl, 0, OxFFFF_FFFF, 0xFFFF_FFFF, temp 32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow < OVh
SPEFSCR
SPEFSCR 
SPEFSCR 
```

For each word element in the accumulator, corresponding even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-68.
If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.
Other registers altered: SPEFSCR ACC


Figure 5-68. Even Form of Vector Half-Word Multiply (evmheusiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmheusianw

User
Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate Negative into Words
evmheusianw rD,rA,rB


```
// high
tempo:31}\leftarrow~r\mp@subsup{A}{0:15}{}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{0:15}{
\mp@subsup{temp}{0:63}{\leftarrowEXTZ (ACC 0:31) - EXTZ (temp 0:31)}
ovh }\leftarrow\mp@subsup{t}{}{\mathrm{ temp}}3
```



```
//low
```



```
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{\mathrm{ ACC 32:63) - EXTZ (temp 0:31 )}}{32}{}
ovl }\leftarrow\mp@subsup{\mathrm{ temp 31}}{}{\prime
rD 32:63}\leftarrow SATURATE (ovl, 0, 0x0000_0000, 0x0000_0000, temp 32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH}< < Ovh
SPEFSCR OV }\leftarrow <ov
SPEFSCR SOVH}\leftarrow < SPEFSCR SOVH | ovh
SPEFSCR 
```

For each word element in the accumulator, corresponding even-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then subtracted from the corresponding word in the accumulator, saturating if underflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-69.
If there is an underflow from the subtraction, the SPEFSCR records overflow and summary overflow bits.
Other registers altered: SPEFSCR ACC


Figure 5-69. Even Form of Vector Half-Word Multiply (evmheusianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhogsmfaa

## evmhogsmfaa

Vector Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate
evmhogsmfaa rD,rA,rB


```
tempo:31}\leftarrowr\mp@subsup{\textrm{rA}}{48:63}{}\mp@subsup{\textrm{X}}{\mathrm{ sf }}{}\mp@subsup{\textrm{rB}}{48:63}{
temp 0:63}\leftarrow EXTS (temp 0:31)
rD0:63}\leftarrow\mp@subsup{\mp@code{ACC}}{0:63}{}+\mp@subsup{\mathrm{ temp 0:63}}{0}{
// update accumulator
ACC 0:63 }\leftarrowr\mp@subsup{r}{0:63}{
```

The corresponding low odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended to 64 bits then added to the contents of the 64-bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-70.

Note: This is a modulo sum. There is no check for overflow and no saturation is performed. An overflow from the 64-bit sum, if one occurs, is not recorded into the SPEFSCR.


Figure 5-70. evmhogsmfaa (Odd Form)

## evmhogsmfan

Vector Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate Negative evmhogsmfan rD,rA,rB



```
temp 0:63}\leftarrow\leftarrow\operatorname{EXTS}(\mp@subsup{t}{emp}{0:31}
rD 0:63}\leftarrow~\mp@subsup{ACC}{0:63 - temp 0:63}{0
// update accumulator
ACC}0:63 \leftarrow rD0:63
```

The corresponding low odd-numbered half-word signed fractional elements in $\mathbf{r} A$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended to 64 bits then subtracted from the contents of the 64 -bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-71.

Note: This is a modulo difference. There is no check for overflow and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-71. evmhogsmfan (Odd Form)

## evmhogsmiaa

| SPE | User |
| :--- | :--- |

evmhogsmiaa
Vector Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer, and Accumulate
evmhogsmiaa rD,rA,rB


The corresponding low odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended to 64 bits then added to the contents of the 64-bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-72.

Note: This is a modulo sum. There is no check for overflow and no saturation is performed. An overflow from the 64-bit sum, if one occurs, is not recorded into the SPEFSCR.


Figure 5-72. evmhogsmiaa (Odd Form)

## evmhogsmian

Vector Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer and Accumulate Negative evmhogsmian rD,rA,rB


```
\mp@subsup{\operatorname{emp}}{0:31}{}\leftarrowr\mp@subsup{A}{48:63}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{rBB}{48:63}{}
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ temp}}{0:31}{}
rD 0:63}\leftarrow & ACC 0:63 - temp 0:63
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

The corresponding low odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is sign-extended to 64 bits then subtracted from the contents of the 64-bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-73.

Note: This is a modulo difference. There is no check for overflow and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-73. evmhogsmian (Odd Form)

## evmhogumiaa

| SPE | User |
| :---: | :---: |

## evmhogumiaa

Vector Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate
evmhogumiaa rD,rA,rB


The corresponding low odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is zero-extended to 64 bits then added to the contents of the 64-bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-74.

Note: This is a modulo sum. There is no check for overflow and no saturation is performed. An overflow from the 64-bit sum, if one occurs, is not recorded into the SPEFSCR.


Figure 5-74. evmhogumiaa (Odd Form)

## evmhogumian

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate Negative
evmhogumian rD,rA,rB


The corresponding low odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is zero-extended to 64 bits then subtracted from the contents of the 64 -bit accumulator, and the result is placed into rD and into the accumulator, as shown in Figure 5-75.

Note: This is a modulo difference. There is no check for overflow and no saturation is performed. Any overflow of the 64-bit difference is not recorded into the SPEFSCR.


Figure 5-75. evmhogumian (Odd Form)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhosmf

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Odd, Signed, Modulo, Fractional (to Accumulator)

| evmhosmf | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmhosmfa | rD,rA,rB | $(A=1)$ |


// high
$r D_{0: 31} \leftarrow\left(r A_{16: 31} \times_{\text {sf }} r B_{16: 31}\right)$
// low
$r D_{32: 63} \leftarrow\left(r_{48: 63} X_{\text {sf }} \quad r B_{48: 63}\right)$
// update accumulator
if $A=1$ then $A C C_{0: 63} \leftarrow r D_{0: 63}$

The corresponding odd-numbered, half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each product is placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-71Figure 5-76.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-76. Vector Multiply Half Words, Odd, Signed, Modulo, Fractional (to Accumulator) (evmhosmf)

## evmhosmfaaw

| SPE | User |
| :---: | :---: |

Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate into Words
evmhosmfaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrow\mp@subsup{rA}{16:31}{}\mp@subsup{\times}{\mathrm{ sf }}{}\mp@subsup{rB}{16:31}{}
rD 0:31}\mp@code{\leftarrow ACC0:31 + temp 0:31
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{t}{}{\prime2mp}0:3
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32 bits of each intermediate product is added to the contents of the corresponding accumulator word and the results are placed into the corresponding $\mathbf{r D}$ words and into the accumulator, as shown in Figure 5-77.
Other registers altered: ACC


Figure 5-77. Odd Form of Vector Half-Word Multiply (evmhosmfaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhosmfanw

| SPE | User |
| :---: | :---: |

Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate Negative into Words
evmhosmfanw rD,rA,rB


```
// high
tempo:31
rD0:31}\leftarrow~\mp@subsup{ACC0:31 - tempo:31}{0}{0
// low
```



```
rD 32:63}\leftarrow\mp@subsup{A}{ACC 32:63 - temp 0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32 bits of each intermediate product is subtracted from the contents of the corresponding accumulator word and the results are placed into the corresponding $\mathbf{r D}$ words and into the accumulator, as shown in Figure 5-78.

Other registers altered: ACC


Figure 5-78. Odd Form of Vector Half-Word Multiply (evmhosmfanw)

## evmhosmi

| SPE | User |
| :---: | :---: |

## evmhosmi

Vector Multiply Half Words, Odd, Signed, Modulo, Integer (to Accumulator)

| evmhosmi | rD,rA,rB | $(\mathrm{A}=0)$ |
| :--- | :--- | :--- |
| evmhosmia | rD,rA,rB | $(\mathrm{A}=1)$ |



```
// high
rD 0:31
// low
rD}\mp@subsup{D}{32:63}{}\leftarrowr\mp@subsup{A}{48:63}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{r|B}{48:63}{
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD 0:63
```

The corresponding odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The two 32-bit products are placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-79.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-79. Vector Multiply Half Words, Odd, Signed, Modulo, Integer (to Accumulator) (evmhosmi)

## evmhosmiaaw

| SPE | User |
| :---: | :---: |

evmhosmiaaw
Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate into Words
evmhosmiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{B}{16:31}{}
rD0:31}\leftarrow\leftarrowACC0:31 + temp 0:31
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{t}{}{\prime2mp}0:3
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate 32-bit product is added to the contents of the corresponding accumulator word and the results are placed into the corresponding $\mathbf{r D}$ words and into the accumulator, as shown in Figure 5-80.
Other registers altered: ACC


Figure 5-80. Odd Form of Vector Half-Word Multiply (evmhosmiaaw)

## evmhosmianw

| SPE | User |
| :---: | :---: |

evmhosmianw
Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate Negative into Words
evmhosmianw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{r\mp@subsup{B}{16:31}{}}{}{\prime}
rD0:31}\leftarrow~ACC0:31 - tempo:31
// low
\mp@subsup{\mathrm{ emp 0:31}}{}{~}\leftarrow\mp@subsup{rA}{48:63 }{ > Si rB4 48:63}
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63 - temp 0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate 32-bit product is subtracted from the contents of the corresponding accumulator word and the results are placed into the corresponding $\mathbf{r D}$ words and into the accumulator, as shown in Figure 5-81.
Other registers altered: ACC


Figure 5-81. Odd Form of Vector Half-Word Multiply (evmhosmianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhossf

Vector Multiply Half Words, Odd, Signed, Saturate, Fractional (to Accumulator)

| evmhossf | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmhossfa | rD,rA,rB | $(A=1)$ |


| 0 |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  | 25 | 26 | 27 |  |  |  | 31 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 0 | 0 | 0 | A | 0 | 0 | 1 | 1 | 1 |

```
// high
```



```
if (rA 16:31 = 0x8000) & (rB 16:31 = 0x8000) then
    rD0:31 \leftarrow0x7FFF_FFFF //saturate
    movh \leftarrow1
else
    rD0:31}\leftarrow\mp@subsup{\mp@code{temp}0:31}{0}{0
    movh \leftarrow0
// low
temp 0:31
if (rA48:63 = 0x8000) & (rB }\mp@subsup{\mp@code{M8:63}}{4}{\prime}=0\times8000) the
    rD 32:63 \leftarrow0x7FFF_FFFF //saturate
    movl \leftarrow1
else
    rD 32:63}\leftarrow\mp@subsup{\mathrm{ temp 0:31}}{0}{
    movl }\leftarrow
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD0:63
// update SPEFSCR
SPEFSCRROVH}\leftarrow\leftarrowmov
SPEFSCR Ov }\leftarrow\mathrm{ movl
SPEFSCR 
SPEFSCR
```

The corresponding odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 32 bits of each product are placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-82. If both inputs are -1.0 , the result saturates to the largest positive signed fraction and the overflow and summary overflow bits are recorded in the SPEFSCR.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: $\quad$ SPEFSCR ACC (If A = 1)


Figure 5-82. Vector Multiply Half Words, Odd, Signed, Saturate, Fractional (to Accumulator) (evmhossf)

Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate into Words
evmhossfaaw rD,rA,rB


```
// high
temp 0:31}\leftarrow~r\mp@subsup{A}{16:31}{}\mp@subsup{x}{\mathrm{ sf }}{}\mp@subsup{rBB}{16:31}{
if (rA 16:31 = 0x8000) & (rB16:31 = 0x8000) then
    tempo:31}\leftarrow0\times7FFF_FFFF //saturate
    movh }\leftarrow
else
    movh \leftarrow0
temp 0:63 \leftarrow EXTS (ACC 0:31) + EXTS (temp 0:31)
ovh }\leftarrow(\mp@subsup{\mathrm{ temp }}{31}{}\oplus\mp@subsup{\mathrm{ temp }}{32}{}
rD0:31 \leftarrowSATURATE (Ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// low
```



```
if (rA 48:63 = 0x8000) & (rB48:63 = 0x8000) then
    tempo:31}\leftarrow0x7FFF_FFFF //saturate
    movl \leftarrow1
else
    movl \leftarrow0
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\operatorname{ACC}}{32:63}{})+\operatorname{EXTS}(\mp@subsup{\mathrm{ temp}}{0:31}{}
ovl }\leftarrow(\mp@subsup{\mathrm{ temp }}{31}{}\oplus\mp@subsup{\mathrm{ temp }}{32}{\prime}
rD 32:63}\leftarrow\mathrm{ SATURATE(Ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCROVH}\leftarrow \leftarrowmovh
SPEFSCR Ov }\leftarrow\mathrm{ movl
SPEFSCR SOVH }\leftarrow\mp@subsup{\mathrm{ SPEFSCR }}{\mathrm{ SOVH }}{|}| ovh | movh
SPEFSCR SOV }\leftarrow\mp@subsup{\mathrm{ SPEFSCR SOV | ovl| movl}}{\mathrm{ SOL}}{
```

The corresponding odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32 -bit product. If both inputs are -1.0 , the result saturates to 0x7FFF_FFFF. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow or underflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-83.

If there is an overflow or underflow from either the multiply or the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-83. Odd Form of Vector Half-Word Multiply (evmhossfaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhossfanw

## Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate Negative into Words

```
evmhossfanw rD,rA,rB
```



```
// high
temp}0:31 \leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ Sf }}{}\mp@subsup{r\mp@subsup{B}{16:31}{}}{}{\prime
if (rA 16:31 = 0x8000) & (rB 16:31 = 0x8000) then
    tempo:31}\leftarrow0x7FFF FFFF //saturate
    movh \leftarrow }
else
    movh}\leftarrow
\mp@subsup{temp}{0:63}{}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC 0:31 }}{0:31}{)}-\operatorname{EXTS}(\mp@subsup{t}{}{\prime}=mp}0:31
ovh \leftarrow(temp 31 \oplus temp 32)
rD 0:31}\leftarrowSATURATE (Ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63
// low
temp 0:31}\leftarrow~\mp@subsup{rA}{48:63}{}\mp@subsup{\times}{\mathrm{ Sf }}{}\mp@subsup{rBB}{48:63}{
if (rA 48:63 = 0x8000) & (rB 48:63 = 0x8000) then
    temp 0:31}\leftarrow00x7FFF_FFFF //saturate
    movl \leftarrow 
else
    movl }\leftarrow
tempo:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC 32:63)}}{0}{})-\operatorname{EXTS}(\mp@subsup{t}{0mp}{0:31}
ovl \leftarrow(\mp@subsup{temp}{31}{}\oplus\mp@subsup{\mathrm{ temp}}{32}{})
rD 32:63}\leftarrow\mathrm{ SATURATE (Ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC}0:63\leftarrowr\mp@subsup{D}{0:63}{0
// update SPEFSCR
SPEFSCR OVH }\leftarrow\mathrm{ movh
SPEFSCR
SPEFSCR 
SPEFSCR SOV 
```

The corresponding odd-numbered half-word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32 -bit product. If both inputs are -1.0 , the result saturates to 0x7FFF_FFFF. Each 32-bit product is then subtracted from the corresponding word in the accumulator, saturating if overflow or underflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-84.

If there is an overflow or underflow from either the multiply or the subtraction, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-84. Odd Form of Vector Half-Word Multiply (evmhossfanw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhossiaaw

| SPE | User |
| :--- | :--- |

## evmhossiaaw

Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate into Words
evmhossiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{rBB}{16:31}{}
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{0:31}{})+\mp@subsup{\operatorname{EXTS}}{(\mp@subsup{t}{}{\prime2mp}}{0:31}
ovh }\leftarrow(\mp@subsup{\mathrm{ temp 31 }}{}{\prime}\mp@subsup{\mathrm{ temp 32}}{2}{\prime}
rD 0:31 \leftarrowSATURATE (Ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// low
temp 0:31}\leftarrow~\mp@subsup{rA}{48:63 (}{~
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{32:63)}{})+\operatorname{EXTS}(\mp@subsup{t}{0mp}{0:31}
ovl \leftarrow(temp 31 \oplus temp 32)
rD 32:63}\leftarrowSATURATE (Ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow\mathrm{ Ovh
SPEFSCR OV }\leftarrow\mathrm{ ovl
SPEFSCR 
SPEFSCR SOV }\leftarrow\mp@subsup{S}{SPEFSCR}{SOV | Ovl
```

The corresponding odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-85.

If there is an overflow or underflow from the addition, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-85. Odd Form of Vector Half-Word Multiply (evmhossiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhossianw

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate Negative into Words
evmhossianw
rD,rA,rB

// high
temp $_{0: 31} \leftarrow r A_{16: 31} \times_{\text {si }} \quad r B_{16: 31}$
temp $_{0: 63} \leftarrow \operatorname{EXTS}\left(\mathrm{ACC}_{0: 31}\right)-\operatorname{EXTS}^{\left(\text {temp }_{0: 31}\right)}$
ovh $\leftarrow\left(\right.$ temp $_{31} \oplus$ temp $\left._{32}\right)$
$r D_{0: 31} \leftarrow \operatorname{SATURATE}\left(\right.$ OVh, $\operatorname{temp}_{31}, 0 \times 8000 \_0000$, 0x7FFF_FFFF, temp $32: 63$ )
// low
temp $_{0: 31} \leftarrow$ rA $_{48: 63} \times_{\text {si }}{r B_{48: 63}}$
temp $_{0: 63} \leftarrow \operatorname{EXTS}\left(\operatorname{ACC}_{32: 63}\right) \quad-\operatorname{EXTS}\left(\right.$ temp $\left._{0: 31}\right)$
ovl $\leftarrow\left(\right.$ temp $_{31} \oplus$ temp $\left._{32}\right)$
$r D_{32}: 63 \leftarrow$ SATURATE (OV1, temp $_{31}, 0 \times 8000 \_0000$, 0x7FFF_FFFF, temp $32: 63$ )
// update accumulator
$\mathrm{ACC}_{0: 63} \leftarrow r \mathrm{D}_{0: 63}$
// update SPEFSCR
SPEFSCR ${ }_{\text {OVH }} \leftarrow$ ovh
SPEFSCR $_{\text {OV }} \leftarrow$ ovl
SPEFSCR $_{\text {SOVH }} \leftarrow$ SPEFSCR $_{\text {SOVH }} \mid$ ovh
SPEFSCR $_{\text {SOV }} \leftarrow$ SPEFSCR $_{\text {SOV }} \mid$ ovl
The corresponding odd-numbered half-word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied, producing a 32-bit product. Each product is subtracted from the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-86.
If there is an overflow or underflow from the subtraction, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC


Figure 5-86. Odd Form of Vector Half-Word Multiply (evmhossianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhoumi

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer (to Accumulator)

| evmhoumi | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmhoumia | rD,rA,rB | $(A=1)$ |



```
// high
rD 0:31}\leftarrow\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{\textrm{rB}}{16:31}{
// low
rD 32:63}\leftarrow~r\mp@subsup{A}{48:63}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rBB}{48:63}{
// update accumulator
if A = 1 then ACC0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The two 32-bit products are placed into the corresponding words of rD, as shown in Figure 5-87.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-87. Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer (to Accumulator) (evmhoumi)

## evmhoumiaaw

| SPE | User |
| :---: | :---: |

Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate into Words
evmhoumiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rB}{16:31}{}
rD0:31}\leftarrow\leftarrowACC0:31 + tempo:31
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{\mathrm{ temp}}{0:31}{
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate product is added to the contents of the corresponding accumulator word. The sums are placed into the corresponding rD and accumulator words, as shown in Figure 5-88.
Other registers altered: ACC


Figure 5-88. Odd Form of Vector Half-Word Multiply (evmhoumiaaw)

## evmhoumianw

| SPE | User |
| :--- | :--- |

## evmhoumianw

Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate Negative into Words
evmhoumianw rD,rA,rB


```
// high
tempo:31}\leftarrowr\mp@subsup{A}{0:15}{}\mp@subsup{\times}{ui}{}\mp@subsup{rBB}{0:15}{
rD0:31 \leftarrow ACC0:31 - temp 0:31
/
/ low
tempo:31}\leftarrowr\mp@subsup{A}{32:47 }{~}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{32:47}{
rD 32:63}\leftarrow~\mp@subsup{ACC}{32:63 - temp 0:31}{
// update accumulator
ACC 0:63}\mp@code{\leftarrow rD 0:63
```

For each word element in the accumulator, the corresponding odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Each intermediate product is subtracted from the contents of the corresponding accumulator word. The results are placed into the corresponding $\mathbf{r D}$ and accumulator words, as shown in Figure 5-89.

Other registers altered: ACC


Figure 5-89. Odd Form of Vector Half-Word Multiply (evmhoumianw)

## evmhousiaaw

| SPE | User |
| :--- | :--- |

Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate into Words
evmhousiaaw rD,rA,rB


```
// high
temp 0:31}\leftarrow~r\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{16:31}{
\mp@subsup{temp}{0:63}{\leftarrowEXXZ (ACC 0:31) + EXTZ (temp 0:31)}\mp@code{EXC}
ovh }\leftarrow\mp@subsup{t}{}{\mathrm{ temp}}3
rD 0:31 \leftarrow SATURATE (ovh, 0, 0xFFFF_FFFF, OxFFFF_FFFF, temp 32:63)
//low
\mp@subsup{emp}{0:31}{}\leftarrowr\mp@subsup{A}{48:63}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rBB}{48:63}{}
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{ACC}{32:63}{})+\operatorname{EXTZ}(\mp@subsup{t}{}{\prime2mp}0:31
ovl }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD 32:63 \leftarrow SATURATE(Ovl, 0, OxFFFF_FFFF, OxFFFF_FFFF, temp 32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH}\leftarrow < ovh
SPEFSCR OV 
SPEFSCR 
SPEFSCR 
```

For each word element in the accumulator, corresponding odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then added to the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-90.

If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.
Other registers altered: SPEFSCR ACC


Figure 5-90. Odd Form of Vector Half-Word Multiply (evmhousiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmhousianw

Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate Negative into Words
evmhousianw rD,rA,rB


```
// high
\mp@subsup{temp}{0:31}{}\leftarrowr\mp@subsup{A}{16:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rB}{16:31}{}
\mp@subsup{temp}{0:63}{\leftarrowE EXTZ (ACC 0:31) - EXTZ (temp 0:31)}
ovh }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD 0:31}\leftarrow\leftarrowSATURATE (Ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp 32:63
//low
```



```
temp 0:63 \leftarrow EXTZ (ACC 32:63) - EXTZ (temp 0:31)
ovl \leftarrow temp 31
rD 32:63 \leftarrow SATURATE(ovl, 0, OxFFFF_FFFF, 0xFFFF_FFFF, temp 32:63)
// update accumulator
ACC 0:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow < OVh
SPEFSCR
SPEFSCR 
SPEFSCR 
```

For each word element in the accumulator, corresponding odd-numbered half-word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 32-bit product. Each 32-bit product is then subtracted from the corresponding word in the accumulator, saturating if overflow occurs, and the result is placed in rD and the accumulator, as shown in Figure 5-91.

If subtraction causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.
Other registers altered: SPEFSCR ACC


Figure 5-91. Odd Form of Vector Half-Word Multiply (evmhousianw)

## evmra

| SPE | User |
| :--- | :--- |

evmra

## Initialize Accumulator

evmra rD,rA


$$
\begin{aligned}
& A C C_{0: 63} \leftarrow r A_{0: 63} \\
& r D_{0: 63} \leftarrow r A_{0: 63}
\end{aligned}
$$

The contents of $\mathbf{r A}$ are written into the accumulator and copied into $\mathbf{r D}$. This is the method for initializing the accumulator, as shown in Figure 5-92.

Other registers altered: ACC


Figure 5-92. Initialize Accumulator (evmra)

## evmwhsmf

SPE

Vector Multiply Word High Signed, Modulo, Fractional (to Accumulator)
evmwhsmf
rD,rA,rB
( $\mathrm{A}=0$ )
evmwhsmfa
rD,rA,rB
( $\mathrm{A}=1$ )


```
// high
tempo:63}\leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ sf }}{}r\mp@subsup{\textrm{B}}{0:31}{
rD0:31}\leftarrow\leftarrow\mp@subsup{temp}{0:31}{
// low
```



```
rD 32:63}\leftarrow\mp@subsup{t}{\mathrm{ temp 0:31}}{
// update accumulator
if A = 1 then ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied and bits $0-31$ of the two products are placed into the two corresponding words of $\mathbf{r D}$, as shown in Figure 5-93.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (if A=1)


Figure 5-93. Vector Multiply Word High Signed, Modulo, Fractional (to Accumulator) (evmwhsmf)

## evmwhsmi

| SPE | User |
| :--- | :--- |

Vector Multiply Word High Signed, Modulo, Integer (to Accumulator)

| evmwhsmi | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmwhsmia | rD,rA,rB | $(A=1)$ |


| 0 |  |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  | 25 | 26 | 27 |  |  |  | 31 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 |  | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 0 | 0 | 1 | A | 0 | 1 | 1 | 0 | 1 | 1 |

```
// high
temp}0:63 \leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{r\mp@subsup{B}{0:31}{}}{0}{
rD 0:31}\leftarrow\mp@subsup{\mathrm{ temp 0:31}}{0}{
// low
temp 0:63}\leftarrowr\mp@subsup{A}{32:63}{}\mp@subsup{\times}{\mathrm{ si }}{}\mp@subsup{rBB}{32:63}{
rD 32:63}\leftarrow\mp@subsup{t}{}{\mathrm{ temp}}0:3
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD 0:63
```

The corresponding word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Bits $0-31$ of the two 64-bit products are placed into the two corresponding words of $\mathbf{r D}$, as shown in Figure 5-94.

If $\mathrm{A}=1$, The result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-94. Vector Multiply Word High Signed, Modulo, Integer (to Accumulator) (evmwhsm)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwhssf

Vector Multiply Word High Signed, Saturate, Fractional (to Accumulator)
evmwhssf
rD,rA,rB
rD,rA,rB
( $\mathrm{A}=1$ )


```
// high
```



```
if (rA 0:31 = 0x8000_0000) & (rB0:31 = 0x8000_0000) then
    rD0:31 \leftarrow0x7FFF_FFFF //saturate
    movh }\leftarrow
else
    rD0:31}\leftarrow\leftarrow\mp@subsup{temp 0:31}{}{0
    movh}\leftarrow
// low
temp 0:63}\leftarrowr\mp@subsup{r}{32:63}{}\mp@subsup{\times}{\mathrm{ sf }}{}\mp@subsup{rB}{32:63}{
if (rA 32:63 = 0x8000_0000) & (rB 32:63 = 0x8000_0000) then
    rD 32:63 \leftarrow0x7FFF_FFFF //saturate
    movl }\leftarrow
else
    rD 32:63}\leftarrow\mp@subsup{\mathrm{ temp 0:31}}{0}{
    movl }\leftarrow
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD0:63
// update SPEFSCR
SPEFSCROvH}\leftarrow\leftarrowmov
SPEFSCR OV 
SPEFSCR SOVH}<<\mp@subsup{\mathrm{ SPEFSCR }}{\mathrm{ SOVH }}{|}|\mathrm{ movh
SPEFSCR
```

The corresponding word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Bits $0-31$ of each product are placed into the corresponding words of $\mathbf{r D}$, as shown in Figure 5-95. If both inputs are -1.0 , the result saturates to the largest positive signed fraction and the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC (If A = 1)


Figure 5-95. Vector Multiply Word High Signed, Saturate, Fractional (to Accumulator) (evmwhssf)
SPE $\quad$ User

Vector Multiply Word High Unsigned, Modulo, Integer (to Accumulator)

| evmwhumi | rD,rA,rB | $(\mathbf{A}=0)$ |
| :--- | :--- | :--- |
| evmwhumia | rD,rA,rB | $(A=1)$ |


| 0 |  |  |  |  | 5 | 6 |  | 10 | 11 |  | 15 | 16 |  | 20 | 21 |  |  |  | 25 | 26 | 27 |  |  |  | 31 | 1 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | 0 | 1 | 0 | 0 |  | rD |  |  | rA |  |  | rB |  | 1 | 0 | 0 | 0 | 1 | A | 0 | 1 | 1 | 0 | 0 |  |

```
// high
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{0:31}{}
rD 0:31}\leftarrow\mp@subsup{\mathrm{ temp 0:31}}{0}{
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{temp}{0:31}{
// update accumulator
if A = 1, ACC 0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. Bits $0-31$ of the two products are placed into the two corresponding words of $\mathbf{r D}$, as shown in Figure 5-96.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-96. Vector Multiply Word High Unsigned, Modulo, Integer (to Accumulator) (evmwhumi)

## evmwlsmiaaw

SPE User

## evmwlsmiaaw

Vector Multiply Word Low Signed, Modulo, Integer and Accumulate in Words
evmwlsmiaaw rD,rA,rB


```
// high
tempo:63}\leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{\textrm{B}}{0:31}{
rD 0:31}\leftarrow\leftarrowACC0:31 + temp 32:63
// low
```



```
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63}{}+\mp@subsup{t}{}{\mathrm{ temp}}32:6
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The least significant 32 bits of each intermediate product is added to the contents of the corresponding accumulator words, and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-97.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: ACC


Figure 5-97. Vector Multiply Word Low Signed, Modulo, Integer and Accumulate in Words (evmwlsmiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwIsmianw

SPE User

## evmwlsmianw

Vector Multiply Word Low Signed, Modulo, Integer and Accumulate Negative in Words
evmwlsmianw rD,rA,rB


```
// high
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{B}{0:31}{}
rD 0:31}\leftarrow\leftarrowACC0:31 - temp 32:63
// low
temp 0:63}\leftarrow\mp@subsup{rAA}{32:63 积 }{\mathrm{ rB }
rD 32:63}\leftarrow\leftarrow\mp@subsup{ACC}{32:63 - temp 32:63}{
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding word elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The least significant 32 bits of each intermediate product is subtracted from the contents of the corresponding accumulator words and the result is placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-98.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: ACC


Figure 5-98. Vector Multiply Word Low Signed, Modulo, Integer and Accumulate Negative in Words (evmwlsmianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwlssiaaw

| SPE | User |
| :--- | :--- |

Vector Multiply Word Low Signed, Saturate, Integer and Accumulate in Words
evmwlssiaaw rD,rA,rB


The corresponding word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 64-bit product. The 32 lsbs of each product are added to the corresponding word in the ACC, saturating if overflow or underflow occurs; the result is placed in $\mathbf{r D}$ and the ACC, as shown in Figure 5-99. If there is overflow or underflow from the addition, overflow and summary overflow bits are recorded in the SPEFSCR.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: SPEFSCR ACC


Figure 5-99. Vector Multiply Word Low Signed, Saturate, Integer and Accumulate in Words (evmwlssiaaw)

## evmwlssianw

| SPE | User |
| :--- | :--- |

## evmwlssianw

Vector Multiply Word Low Signed, Saturate, Integer and Accumulate Negative in Words
evmwlssianw rD,rA,rB


```
// high
temp}0:63 \leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ si }}{}r\mp@subsup{r}{0:31}{
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC }}{0:31}{0}) - EXTS (temp 32:63) 
ovh }\leftarrow(\mp@subsup{\mathrm{ temp 31 }}{}{\prime}\mp@subsup{\mathrm{ temp 32}}{2}{\prime}
rD 0:31 \leftarrowSATURATE (Ovh, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// low
temp 0:63}\leftarrowr\mp@subsup{A}{32:63 (}{~
temp 0:63}\leftarrow\operatorname{EXTS}(\mp@subsup{\mathrm{ ACC 32:63) - EXTS (temp 32:63)}}{32}{}
ovl \leftarrow }\leftarrow(\mp@subsup{\mathrm{ temp}}{31}{}\oplus\mp@subsup{\mathrm{ temp 32 )}}{3}{\prime
rD 32:63}\leftarrowSATURATE (Ovl, temp 31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow\mathrm{ ovh
SPEFSCR OV }\leftarrow\mathrm{ ovl
SPEFSCR 
SPEFSCR 
```

The corresponding word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 64-bit product. The 32 lsbs of each product are subtracted from the corresponding ACC word, saturating if overflow or underflow occurs, and the result is placed in rD and the ACC, as shown in Figure 5-100. If addition causes overflow or underflow, overflow and summary overflow SPEFSCR bits are recorded.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: SPEFSCR ACC


Figure 5-100. Vector Multiply Word Low Signed, Saturate, Integer and Accumulate Negative in Words (evmwlssianw

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwlumi

| SPE | User |
| :---: | :---: |

Vector Multiply Word Low Unsigned, Modulo, Integer

| evmwlumi | rD,rA,rB | $(\mathrm{A}=0)$ |
| :--- | :--- | :--- |
| evmwlumia | rD,rA,rB | $(\mathrm{A}=1)$ |



```
// high
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{0:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}r\mp@subsup{B}{0:31}{}
rD}0:31 \leftarrow \leftarrowemp 32:63
// low
temp 0:63}\leftarrowr\mp@subsup{A}{32:63}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{r~B}{32:63}{
rD 32:63}\leftarrow\mp@subsup{t}{}{\mathrm{ temp}}32:6
// update accumulator
If A = 1 then ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

The corresponding word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The least significant 32 bits of each product are placed into the two corresponding words of $\mathbf{r D}$, as shown in Figure 5-101.

Note: The least significant 32 bits of the product are independent of whether the word elements in $\mathbf{r A}$ and $\mathbf{r B}$ are treated as signed or unsigned 32-bit integers.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)
Note that evmwlumi and evmwlumia can be used for signed or unsigned integers.


Figure 5-101. Vector Multiply Word Low Unsigned, Modulo, Integer (evmwlumi)

## evmwlumiaaw

| SPE | User |
| :---: | :---: |

## evmwlumiaaw

Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate in Words
evmwlumiaaw rD,rA,rB


```
// high
temp}0:63 \leftarrowr\mp@subsup{A}{0:31 }{*}\mp@subsup{X}{ui}{}r\mp@subsup{B}{0:31}{
rD 0:31}\leftarrow\leftarrowACC0:31 + temp 32:63
// low
```



```
rD 32:63}\leftarrow\mp@subsup{\textrm{ACC}}{32:63}{}+\mp@subsup{\mathrm{ temp}}{32:63}{
// update accumulator
ACC 0:63}\leftarrowr\mp@subsup{D}{0:63}{
```

For each word element in the accumulator, the corresponding word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The least significant 32 bits of each product are added to the contents of the corresponding accumulator word and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-102.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: ACC


Figure 5-102. Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate in Words (evmwlumiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwlumianw

| SPE | User |
| :---: | :---: |

Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate Negative in Words
evmwlumianw rD,rA,rB


```
// high
tempo:63}\leftarrow\mp@subsup{rA}{0:31}{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rBB}{0:31}{
rD0:31}\leftarrow\leftarrow\mp@subsup{ACC}{0:31 - temp 32:63}{
// low
temp 0:63}\leftarrowr\mp@subsup{A}{32:63 }{}\mp@subsup{\times}{\mathrm{ ui }}{}\mp@subsup{rBB}{32:63}{
rD 32:63}\leftarrow~\mp@subsup{ACC}{32:63 - temp 32:63}{
// update accumulator
ACC
```

For each word element in the accumulator, the corresponding word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The least significant 32 bits of each product are subtracted from the contents of the corresponding accumulator word and the result is placed into $\mathbf{r D}$ and the ACC, as shown in Figure 5-103.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: ACC


Figure 5-103. Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate Negative in Words (evmwlumianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwlusiaaw

SPE $\quad$ User

Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate in Words
evmwlusiaaw rD,rA,rB


```
// high
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{0:31}{}}\mp@subsup{\times}{ui}{}\mp@subsup{rBB}{0:31}{
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{\mathrm{ ACC }}{0:31}{})+\mp@subsup{\operatorname{EXTZ}}{(temp}{32:63}
ovh }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD 0:31}\leftarrow SATURATE (Ovh, 0, 0xFFFF_FFFF, OxFFFF_FFFF, temp 32:63
//low
temp 0:63}\leftarrowr\mp@subsup{A}{32:63}{}\mp@subsup{X}{\mathrm{ ui }}{}r\mp@subsup{B}{32:63}{
temp 0:63}\leftarrowEXTZ (ACC 32:63) + EXTZ (temp 32:63)
ovl }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD 32:63}\leftarrow SATURATE (Ovl, 0, OxFFFF_FFFF, 0xFFFF_FFFF, temp 32:63
// update accumulator
ACC 0:63}\leftarrow rD 0:63
// update SPEFSCR
SPEFSCR 
SPEFSCR OV }\leftarrow < ov
SPEFSCR 
SPEFSCR 
```

For each word element in the ACC , corresponding word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied, producing a 64 -bit product. The 32 lsbs of each product are added to the corresponding ACC word, saturating if overflow occurs; the result is placed in rD and the ACC, as shown in Figure 5-104. If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: SPEFSCR ACC


Figure 5-104. Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate in Words (evmwlusiaaw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate Negative in Words
evmwlusianw rD,rA,rB


```
// high
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{0:31}{}}\mp@subsup{\times}{ui}{}\mp@subsup{rBB}{0:31}{
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{\mathrm{ ACC 0:31 }}{0:31}{})-\mp@subsup{\operatorname{EXTZ}}{(temp}{32:63}
ovh }\leftarrow\mp@subsup{\mathrm{ temp 31}}{3}{
rD D:31}\leftarrow \leftarrowATURATE (ovh, 0, 0x0000_0000, 0x0000_0000, temp 32:63
//low
```



```
temp 0:63}\leftarrow\operatorname{EXTZ}(\mp@subsup{\textrm{ACC}}{32:63}{0})-\mp@subsup{\textrm{EXTZ}}{(temp}{32:63}
ovl }\leftarrow\mp@subsup{t}{}{\mathrm{ temp}}3
rD 32:63}\leftarrow SATURATE (ovl, 0, 0x0000_0000, 0x0000_0000, temp 32:63)
// update accumulator
ACC 0:63}\leftarrow rD 0:63
// update SPEFSCR
SPEFSCR OVH }\leftarrow Ov
SPEFSCR OV }\leftarrow Ov
SPEFSCR 
SPEFSCR 
```

For each ACC word element, corresponding word elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 64-bit product. The 32 lsbs of each product are subtracted from corresponding ACC words, saturating if underflow occurs; the result is placed in $\mathbf{r D}$ and the ACC, as shown in Figure 5-105. If there is an underflow from the subtraction, the overflow and summary overflow bits are recorded in the SPEFSCR.

## NOTE

Care should be taken if the intermediate product cannot be represented in 32 bits as some implementations produce an undefined final result. Status bits are set that indicate that such an overflow occurred.

Other registers altered: SPEFSCR ACC


Figure 5-105. Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate Negative in Words (evmwlusianw)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwsmf

| SPE | User |
| :---: | :---: |

Vector Multiply Word Signed, Modulo, Fractional (to Accumulator)
evmwsmf
rD,rA,rB
evmwsmfa
rD,rA,rB



```
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD0:63
```

( $\mathrm{A}=0$ )
( $\mathrm{A}=1$ )

The corresponding low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The product is placed into rD, as shown in Figure 5-106.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: ACC (If A=1)


Figure 5-106. Vector Multiply Word Signed, Modulo, Fractional (to Accumulator) (evmwsmf)

## evmwsmfaa

Vector Multiply Word Signed, Modulo, Fractional and Accumulate
evmwsmfaa
rD,rA,rB



```
rD 0:63}\leftarrow\leftarrow\mp@subsup{\textrm{ACC}}{0:63}{}+\mp@subsup{\mathrm{ temp}}{0:63}{
// update accumulator
ACC}0:63~r\mp@subsup{D}{0:63}{
```

The corresponding low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is added to the contents of the 64-bit accumulator and the result is placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-107.

Other registers altered: ACC


Figure 5-107. Vector Multiply Word Signed, Modulo, Fractional and Accumulate (evmwsmfaa)

## evmwsmfan

Vector Multiply Word Signed, Modulo, Fractional and Accumulate Negative
evmwsmfan
rD,rA,rB


The corresponding low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is subtracted from the contents of the accumulator and the result is placed in $\mathbf{r D}$ and the accumulator, as shown in Figure 5-108.

Other registers altered: ACC


Figure 5-108. Vector Multiply Word Signed, Modulo, Fractional and Accumulate Negative (evmwsmfan)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwsmi

Vector Multiply Word Signed, Modulo, Integer (to Accumulator)
evmwsmi
rD,rA,rB
evmwsmia
rD,rA,rB


```
rD 0:63}\mp@code{\leftarrowrA 32:63 }\mp@subsup{X}{\mathrm{ si }}{}r\mp@subsup{B}{32:63}{
// update accumulator
if A = 1 then ACC 0:63 \leftarrowrD 0:63
```

( $\mathrm{A}=0$ )
( $\mathrm{A}=1$ )

The low word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The product is placed into $\mathbf{r D}$.
If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator., as shown in Figure 5-109.
Other registers altered: ACC (If A=1)


Figure 5-109. Vector Multiply Word Signed, Modulo, Integer (to Accumulator) (evmwsmi)

## evmwsmiaa

Vector Multiply Word Signed, Modulo, Integer and Accumulate
evmwsmiaa
rD,rA,rB


```
temp 0:63}\mp@code{\leftarrowrA 32:63 }\mp@subsup{~}{\mathrm{ si }}{}\mp@subsup{rBB}{32:63}{
rD 0:63}\leftarrow\leftarrow\mp@subsup{ACC}{0:63}{}+\mp@subsup{\mathrm{ temp}}{0:63}{
// update accumulator
ACC}0:63~r\mp@subsup{D}{0:63}{
```

The low word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is added to the contents of the 64-bit accumulator and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-110.

Other registers altered: ACC


Figure 5-110. Vector Multiply Word Signed, Modulo, Integer and Accumulate (evmwsmiaa)

## evmwsmian

| SPE | User |
| :--- | :--- |

## evmwsmian

Vector Multiply Word Signed, Modulo, Integer and Accumulate Negative
evmwsmian
rD,rA,rB


```
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{32:63}{}\mp@subsup{X}{\mathrm{ si }}{}\mp@subsup{r\mp@subsup{B}{32:63}{}}{}{\prime2}
rD 0:63}\leftarrow\leftarrow\mp@subsup{ACC}{0:63 - temp 0:63}{
// update accumulator
ACC}0:63~r\mp@subsup{D}{0:63}{
```

The low word signed integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is subtracted from the contents of the 64-bit accumulator and the result is placed into $\mathbf{r D}$ and the accumulator, as shown in Figure 5-111.

Other registers altered: ACC


Figure 5-111. Vector Multiply Word Signed, Modulo, Integer and Accumulate Negative (evmwsmian)

## evmwssf

| SPE | User |
| :---: | :---: |

evmwssf
Vector Multiply Word Signed, Saturate, Fractional (to Accumulator)

| evmwssf | rD,rA,rB | $(\mathrm{A}=0)$ |
| :--- | :--- | :--- |
| evmwssfa | rD,rA,rB | $(\mathrm{A}=1)$ |


temp $_{0: 63} \leftarrow r A_{32: 63} \quad \times_{\text {sf }} \quad r B_{32: 63}$
if $\left(r A_{32: 63}=0 \times 8000 \_0000\right) \&\left(r B_{32: 63}=0 \times 8000 \_0000\right)$ then
rD $0: 63 \leftarrow 0 \times 7$ FFF_FFFF_FFFF_FFFF / /saturat $\bar{e}$
mov $\leftarrow 1$
else
$r D_{0: 63} \leftarrow$ temp $_{0: 63}$
mov $\leftarrow 0$
// update accumulator
if $A=1$ then $A C C_{0: 63} \leftarrow r D_{0: 63}$
// update SPEFSCR
SPEFSCR $_{\text {OVH }} \leftarrow 0$
$\mathrm{SPEFSCR}_{\mathrm{OV}} \leftarrow \mathrm{mov}$
SPEFSCR $_{\text {SOV }} \leftarrow$ SPEFSCR $_{\text {SOV }} \mid \mathrm{mov}$

The low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The 64 bit product is placed into $\mathbf{r D}$, as shown in Figure 5-112. If both inputs are -1.0 , the result saturates to the largest positive signed fraction and the overflow and summary overflow bits are recorded in the SPEFSCR.

The architecture specifies that if the final result cannot be represented in 64 bits, SPEFSCR[OV] should be set (along with the SOV bit, if it is not already set).

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: $\quad$ SPEFSCR ACC (If A = 1)


Figure 5-112. Vector Multiply Word Signed, Saturate, Fractional (to Accumulator) (evmwssf)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwssfaa

| SPE | User |
| :--- | :--- |

## evmwssfaa

Vector Multiply Word Signed, Saturate, Fractional and Accumulate
evmwssfaa
rD,rA,rB


```
\mp@subsup{temp}{0:63}{}\leftarrowr\mp@subsup{A}{32:63}{}\mp@subsup{X}{\mathrm{ sf }}{}\mp@subsup{rBB}{32:63}{}
if (rA 32:63 = 0x8000_0000) & (rB 32:63 = 0x8000_0000) then
    tempo:63 \leftarrow0x7FFF_FFFF_FFFF_FFFF //saturate
    mov \leftarrow < 
else
    mov }\leftarrow
temp 0:64}\leftarrow\mathrm{ EXTS (ACC 0:63) + EXTS (temp 0:63)
ov }\leftarrow(\mp@subsup{t}{}{\prime
rD 0:63}\mp@code{\leftarrowtemp 1:64 )
// update accumulator
ACCO:63}\leftarrow~r\mp@subsup{D}{0:63}{
// update SPEFSCR
SPEFSCR OVH }\leftarrow
SPEFSCROv }\leftarrow\mathrm{ mov
SPEFSCR SOV }\leftarrow\mp@subsup{\mathrm{ SPEFSCR }}{\mathrm{ SOv | | ov | mov}}{
```

The low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 64-bit product. If both inputs are -1.0 , the product saturates to the largest positive signed fraction. The 64-bit product is added to the ACC and the result is placed in $\mathbf{r D}$ and the ACC, as shown in Figure 5-113.
If there is an overflow from either the multiply or the addition, the SPEFSCR overflow and summary overflow bits are recorded.

Note: There is no saturation on the addition with the accumulator.
Other registers altered: SPEFSCR ACC


Figure 5-113. Vector Multiply Word Signed, Saturate, Fractional, and Accumulate (evmwssfaa)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwssfan

| SPE | User |
| :--- | :--- |

## evmwssfan

Vector Multiply Word Signed, Saturate, Fractional and Accumulate Negative
evmwssfan
rD,rA,rB


The low word signed fractional elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied producing a 64 -bit product. If both inputs are -1.0 , the product saturates to the largest positive signed fraction. The 64-bit product is subtracted from the ACC and the result is placed in $\mathbf{r D}$ and the ACC, as shown in Figure 5-114.

If there is an overflow from either the multiply or the addition, the SPEFSCR overflow and summary overflow bits are recorded.

Note: There is no saturation on the subtraction with the accumulator.
Other registers altered: SPEFSCR ACC


Figure 5-114. Vector Multiply Word Signed, Saturate, Fractional and Accumulate Negative (evmwssfan)

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## evmwumi

| SPE | User |
| :---: | :---: |

evmwumi
Vector Multiply Word Unsigned, Modulo, Integer (to Accumulator)
evmwumi
rD,rA,rB
( $\mathrm{A}=0$ )
evmwumia
rD,rA,rB
( $\mathrm{A}=1$ )


```
rD 0:63}\mp@code{\leftarrowr\mp@subsup{A}{32:63 }{* Xi rB}
// update accumulator
if A = 1 then ACC 0:63 }\leftarrowr\mp@subsup{D}{0:63}{
```

The low word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied to form a 64-bit product that is placed into rD, as shown in Figure 5-115.

If $\mathrm{A}=1$, the result in $\mathbf{r D}$ is also placed into the accumulator.
Other registers altered: $\mathrm{ACC}($ If $\mathrm{A}=1)$


Figure 5-115. Vector Multiply Word Unsigned, Modulo, Integer (to Accumulator) (evmwumi)

## evmwumiaa

| SPE | User |
| :---: | :---: |

## evmwumiaa

Vector Multiply Word Unsigned, Modulo, Integer and Accumulate
evmwumiaa
rD,rA,rB


The low word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is added to the contents of the 64-bit accumulator, and the resulting value is placed into the accumulator and into rD, as shown in Figure 5-116.

Other registers altered: ACC


Figure 5-116. Vector Multiply Word Unsigned, Modulo, Integer and Accumulate (evmwumiaa)

## evmwumian

| SPE | User |
| :--- | :--- |

## evmwumian

Vector Multiply Word Unsigned, Modulo, Integer and Accumulate Negative
evmwumian rD,rA,rB


The low word unsigned integer elements in $\mathbf{r A}$ and $\mathbf{r B}$ are multiplied. The intermediate product is subtracted from the contents of the 64-bit accumulator, and the resulting value is placed into the accumulator and into rD, as shown in Figure 5-117.

Other registers altered: ACC


Figure 5-117. Vector Multiply Word Unsigned, Modulo, Integer and Accumulate Negative (evmwumian)

## evnand

SPE $\quad$ User
evnand
Vector NAND
evnand
rD,rA,rB


```
rD}0:31~\neg(r\mp@subsup{A}{0:31 & rB}{0:31
rD 32:63}\leftarrow\neg(r\mp@subsup{A}{32:63 & rB 32:63) // Bitwise NAND}{
```

Corresponding word elements of $\mathbf{r A}$ and $\mathbf{r B}$ are bitwise NANDed. The result is placed in the corresponding element of rD, as shown in Figure 5-118.


Figure 5-118. Vector NAND (evnand)

## evneg

| SPE | User |
| :--- | :--- |

## Vector Negate

evneg
rD,rA


$$
\begin{aligned}
& \mathrm{rD}_{0: 31} \leftarrow \operatorname{NEG}\left(r A_{0: 31}\right) \\
& r D_{32}: 63 \\
& \leftarrow \operatorname{NEG}\left(r A_{32: 63}\right)
\end{aligned}
$$

The negative of each element of $\mathbf{r A}$ is placed in $\mathbf{r D}$, as shown in Figure 5-119. The negative of 0x8000_0000 (most negative number) returns 0x8000_0000. No overflow is detected.


Figure 5-119. Vector Negate (evneg)

## evnor

| SPE | User |
| :--- | :--- |

evnor
Vector NOR
evnor
rD,rA,rB

$r D_{0: 31} \leftarrow \neg\left(r A_{0: 31} \mid \underset{0: 31}{ }\right) / /$ Bitwise NOR
$r D_{32: 63} \leftarrow \neg\left(r A_{32: 63} \mid r B_{32: 63}\right) / /$ Bitwise NOR
Each element of $\mathbf{r A}$ and $\mathbf{r B}$ is bitwise NORed. The result is placed in the corresponding element of $\mathbf{r D}$, as shown in Figure 5-120.

Note: Use evnand or evnor for evnot.


Figure 5-120. Vector NOR (evnor)
Simplified mnemonic: evnot rD,rA performs a complement register

> evnot rD,rA
equivalent to
evnor rD,rA,rA

## evor

| SPE | User |
| :--- | :--- |

Vector OR
evor
rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow r A_{0: 31} \mid r B_{0: 31} / / \text { Bitwise OR } \\
& r D_{32: 63} \leftarrow r A_{32: 63} \mid r B_{32: 63} / / \text { Bitwise OR }
\end{aligned}
$$

Each element of $\mathbf{r A}$ and $\mathbf{r B}$ is bitwise ORed. The result is placed in the corresponding element of $\mathbf{r D}$, as shown in Figure 5-121.


Figure 5-121. Vector OR (evor)
Simplified mnemonic: evmr rD,rA handles moving of the full 64-bit SPE register.
evmr rD,rA
equivalent to
evor rD,rA,rA

## evorc

| SPE | User |
| :--- | :--- |

evorc
Vector OR with Complement
evorc
rD,rA,rB


$$
\begin{aligned}
& r D_{0: 31} \leftarrow r A_{0: 31} \mid\left(\neg r B_{0: 31}\right) / / \text { Bitwise ORC } \\
& r D_{32: 63} \leftarrow r A_{32: 63} \mid\left(\neg r B_{32: 63}\right) \text { // Bitwise ORC }
\end{aligned}
$$

Each element of $\mathbf{r A}$ is bitwise ORed with the complement of $\mathbf{r B}$. The result is placed in the corresponding element of rD, as shown in Figure 5-122.


Figure 5-122. Vector OR with Complement (evorc)

## evrlw

| SPE | User |
| :--- | :--- |

## Vector Rotate Left Word

evrlw rD,rA,rB


```
nh}\leftarrowr\mp@subsup{\textrm{B}}{27:31}{
nl}\leftarrow\mp@subsup{rgB}{59:63}{
rD 0:31}\leftarrow~ROTL (rA 0:31, nh)
rD 32:63}\leftarrow~\operatorname{ROTL}(r\mp@subsup{A}{32:63, nl)}{n
```

Each of the high and low elements of $\mathbf{r A}$ is rotated left by an amount specified in $\mathbf{r B}$. The result is placed into $\mathbf{r D}$, as shown in Figure 5-123. Rotate values for each element of $\mathbf{r A}$ are found in bit positions $\mathbf{r B}[27-31]$ and $\mathbf{r B}[59-63]$.


Figure 5-123. Vector Rotate Left Word (evrlw)

## evrlwi

| SPE | User |
| :--- | :--- |

evrlwi
Vector Rotate Left Word Immediate
evrlwi rD,rA,UIMM

$\mathrm{n} \leftarrow \mathrm{UIMM}$
$r D_{0: 31} \leftarrow \operatorname{ROTL}\left(r A_{0: 31}, \mathrm{n}\right)$
$r D_{32: 63} \leftarrow \operatorname{ROTL}\left(r A_{32: 63}, \mathrm{n}\right)$
Both the high and low elements of $\mathbf{r A}$ are rotated left by an amount specified by a 5-bit immediate value, as shown in Figure 5-124.


Figure 5-124. Vector Rotate Left Word Immediate (evrlwi)

## evrndw

| SPE | User |
| :---: | :---: |

evrndw
Vector Round Word
evrndw rD,rA

$r D_{0: 31} \leftarrow\left(r A_{0: 31}+0 \times 00008000\right) \& 0 x F F F F 0000 / /$ Modulo sum
$r D_{32: 63} \leftarrow\left(r A_{32: 63}+0 x 00008000\right) \& 0 x F F F F 0000 / /$ Modulo sum
The 32-bit elements of $\mathbf{r A}$ are rounded into 16 bits. The result is placed into $\mathbf{r D}$, as shown in Figure 5-125. The resulting 16 bits are placed in the most significant 16 bits of each element of $\mathbf{r D}$, zeroing out the low order 16 bits of each element.


Figure 5-125. Vector Round Word (evrndw)

## evsel

| SPE | User |
| :---: | :---: |

## Vector Select

evsel
rD,rA,rB,crS


$$
\begin{aligned}
& \text { ch } \leftarrow \mathrm{CR}_{\mathrm{CrS} * 4} \\
& \text { cl } \leftarrow \mathrm{CR}_{\mathrm{CrS} * 4+1} \\
& \text { if }(\mathrm{ch}=1) \text { then } r D_{0: 31} \leftarrow r A_{0: 31} \\
& \text { else } r D_{0: 31} \leftarrow r B_{0: 31} \\
& \text { if }(c l=1) \text { then } r D_{32: 63} \leftarrow r A_{32: 63} \\
& \text { else } r D_{32: 63} \leftarrow r B_{32: 63}
\end{aligned}
$$

If the most significant bit in the crS field of CR is set, the high-order element of $\mathbf{r A}$ is placed in the high-order element of $\mathbf{r D}$; otherwise, the high-order element of $\mathbf{r B}$ is placed into the high-order element of $\mathbf{r D}$. If the next most significant bit in the crS field of CR is set, the low-order element of $\mathbf{r A}$ is placed in the low-order element of $\mathbf{r D}$, otherwise, the low-order element of $\mathbf{r B}$ is placed into the low-order element of rD. This is shown in Figure 5-126.


Figure 5-126. Vector Select (evsel)

## evslw

| SPE | User |
| :---: | :---: |

Vector Shift Left Word
evslw rD,rA,rB


$$
\begin{aligned}
& \mathrm{nh} \leftarrow r \mathrm{~B}_{26: 31} \\
& \mathrm{nl} \leftarrow r \mathrm{~B}_{58: 63} \\
& r D_{0: 31} \leftarrow \mathrm{SL}\left(r \mathrm{~A}_{0: 31}, \mathrm{nh}\right) \\
& r D_{32: 63} \leftarrow \mathrm{SL}\left(r \mathrm{~A}_{32: 63, \mathrm{nl})}\right.
\end{aligned}
$$

Each of the high and low elements of $\mathbf{r A}$ are shifted left by an amount specified in $\mathbf{r B}$. The result is placed into $\mathbf{r D}$, as shown in Figure 5-127. The separate shift amounts for each element are specified by 6 bits in rB that lie in bit positions 26-31 and 58-63.
Shift amounts from 32 to 63 give a zero result.


Figure 5-127. Vector Shift Left Word (evsiw)

## evslwi

| SPE | User |
| :--- | :--- |

evsIwi
Vector Shift Left Word Immediate
evslwi rD,rA,UIMM

$\mathrm{n} \leftarrow$ UIMM
$r D_{0: 31} \leftarrow S L\left(r A_{0: 31}, n\right)$
$r D_{32: 63} \leftarrow S L\left(r A_{32: 63}, n\right)$
Both high and low elements of $\mathbf{r A}$ are shifted left by the 5-bit UIMM value and the results are placed in rD, as shown in Figure 5-128.


Figure 5-128. Vector Shift Left Word Immediate (evslwi)

## evsplatfi

Vector Splat Fractional Immediate
evsplatfi rD,SIMM

$r D_{0: 31} \leftarrow$ SIMM $\left|\left\lvert\, \begin{array}{c}270 \\ r D_{32: 63}\end{array}\right.\right.$ SIMM $\left.^{27}\right|{ }^{27} 0$
The 5-bit immediate value is padded with trailing zeros and placed in both elements of $\mathbf{r D}$, as shown in Figure 5-129. The SIMM ends up in bit positions rD[0-4] and rD[32-36].


Figure 5-129. Vector Splat Fractional Immediate (evsplatfi)

## evsplati

| SPE | User |
| :--- | :--- |

Vector Splat Immediate
evsplati rD,SIMM

$r D_{0: 31} \leftarrow \operatorname{EXTS}(S I M M)$
$r D_{32: 63} \leftarrow \operatorname{EXTS}($ SIMM $)$
The 5-bit immediate value is sign extended and placed in both elements of $\mathbf{r D}$, as shown in Figure 5-130.


Figure 5-130. evsplati Sign Extend

## evsrwis

|  | User |
| :--- | :--- |

Vector Shift Right Word Immediate Signed
evsrwis rD,rA,UIMM

$\mathrm{n} \leftarrow \mathrm{UIMM}$
$r D_{0: 31} \leftarrow \operatorname{EXTS}\left(r A_{0: 31-n}\right)$
$r D_{32: 63} \leftarrow \operatorname{EXTS}\left(r A_{32: 63-n}\right)$
Both high and low elements of $\mathbf{r A}$ are shifted right by the 5-bit UIMM value, as shown in Figure 5-131. Bits in the most significant positions vacated by the shift are filled with a copy of the sign bit.


Figure 5-131. Vector Shift Right Word Immediate Signed (evsrwis)

## evsrwiu

| SPE | User |
| :---: | :---: |

evsrwiu
Vector Shift Right Word Immediate Unsigned
evsrwiu rD,rA,UIMM

$\mathrm{n} \leftarrow \mathrm{UIMM}$
$r D_{0: 31} \leftarrow \operatorname{EXTZ}\left(r A_{0: 31-n}\right)$
$r D_{32: 63} \leftarrow \operatorname{EXTZ}\left(r A_{32: 63-n}\right)$
Both high and low elements of $\mathbf{r A}$ are shifted right by the 5-bit UIMM value; 0 bits are shifted in to the most significant position, as shown in Figure 5-132. Bits in the most significant positions vacated by the shift are filled with a zero bit.


Figure 5-132. Vector Shift Right Word Immediate Unsigned (evsrwiu)

## evsrws

| SPE | User |
| :---: | :---: |

## evsrws

Vector Shift Right Word Signed
evsrws rD,rA,rB


$$
\begin{aligned}
& n h \leftarrow r B_{26: 31} \\
& n l \leftarrow r B_{58: 63} \\
& r D_{0: 31} \leftarrow \operatorname{EXTS}\left(r A_{0: 31-n h}\right) \\
& r D_{32}: 63 \leftarrow \operatorname{EXTS}\left(r A_{32: 63-n l}\right)
\end{aligned}
$$

Both the high and low elements of $\mathbf{r A}$ are shifted right by an amount specified in $\mathbf{r B}$. The result is placed into rD, as shown in Figure 5-133. The separate shift amounts for each element are specified by 6 bits in rB that lie in bit positions 26-31 and 58-63. The sign bits are shifted in to the most significant position.
Shift amounts from 32 to 63 give a result of 32 sign bits.


Figure 5-133. Vector Shift Right Word Signed (evsrws)

## evsrwu

| SPE | User |
| :--- | :--- |

## Vector Shift Right Word Unsigned

evsrwu rD,rA,rB


$$
\begin{aligned}
& n h \leftarrow r B_{26: 31} \\
& n l \leftarrow r B_{58: 63} \\
& r D_{0: 31} \leftarrow \operatorname{EXTZ}\left(r A_{0: 31-n h}\right) \\
& r D_{32}: 63 \leftarrow \operatorname{EXTZ}\left(r A_{32: 63-n l}\right)
\end{aligned}
$$

Both the high and low elements of $\mathbf{r A}$ are shifted right by an amount specified in $\mathbf{r B}$. The result is placed into rD, as shown in Figure 5-134. The separate shift amounts for each element are specified by 6 bits in $\mathbf{r B}$ that lie in bit positions 26-31 and 58-63. Zero bits are shifted in to the most significant position.

Shift amounts from 32 to 63 give a zero result.


Figure 5-134. Vector Shift Right Word Unsigned (evsrwu)

| SPE, SPE FV, SPE FD | User |
| :---: | :---: |

Vector Store Double of Double

The contents of $\mathbf{r S}$ are stored as a double word in storage addressed by EA, as shown in Figure 5-135.
Figure 5-135 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |



Figure 5-135. evstdd Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstddx

SPE, SPE FV, SPE FD $\quad$ User

Vector Store Double of Double Indexed
evstddx rS,rA,rB

if $(r A=0)$ then $b \leftarrow 0$
else $\mathrm{b} \leftarrow(r \mathrm{~A})$
$\mathrm{EA} \leftarrow \mathrm{b}+(r \mathrm{~B})$
$\operatorname{MEM}(E A, 8) \leftarrow R S_{0: 63}$
The contents of $\mathbf{r S}$ are stored as a double word in storage addressed by EA.
Figure 5-136 shows how bytes are stored in memory as determined by the endian mode.


Figure 5-136. evstddx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstdh

| SPE | User |
| :--- | :--- |

Vector Store Double of Four Half Words

```
evstdh rS,d(rA)
```



```
if (rA \(=0\) ) then \(\mathrm{b} \leftarrow 0\)
else \(\mathrm{b} \leftarrow(r \mathrm{~A})\)
EA \(\leftarrow \mathrm{b}+\) EXTZ (UIMM*8)
\(\operatorname{MEM}(E A, 2) \leftarrow R S_{0: 15}\)
\(\operatorname{MEM}(\mathrm{EA}+2,2) \leftarrow \mathrm{RS}_{16: 31}\)
\(\operatorname{MEM}(\mathrm{EA}+4,2) \leftarrow \mathrm{RS}_{32: 47}\)
\(\operatorname{MEM}(\mathrm{EA}+6,2) \leftarrow \mathrm{RS}_{48: 63}\)
```

The contents of $\mathbf{r S}$ are stored as four half words in storage addressed by EA.
Figure 5-137 shows how bytes are stored in memory as determined by the endian mode.


Figure 5-137. evstdh Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstdhx

| SPE | User |
| :--- | :--- |

Vector Store Double of Four Half Words Indexed
evstdhx rS,rA,rB

if $(r A=0)$ then $b \leftarrow 0$
else $\mathrm{b} \leftarrow(r \mathrm{~A})$
$\mathrm{EA} \leftarrow \mathrm{b}+(r \mathrm{~B})$
$\operatorname{MEM}(E A, 2) \leftarrow R S_{0: 15}$
$\operatorname{MEM}(\mathrm{EA}+2,2) \leftarrow \mathrm{RS}_{16: 31}$
$\operatorname{MEM}(\mathrm{EA}+4,2) \leftarrow \mathrm{RS}_{32: 47}$
$\operatorname{MEM}(\mathrm{EA}+6,2) \leftarrow \mathrm{RS}_{48: 63}$
The contents of $\mathbf{r S}$ are stored as four half words in storage addressed by EA.
Figure 5-138 shows how bytes are stored in memory as determined by the endian mode.


Figure 5-138. evstdhx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstdw

| SPE | User |
| :--- | :--- |

Vector Store Double of Two Words
evstdw rS,d(rA)


```
if (rA = 0) then b }\leftarrow
```

else $\mathrm{b} \leftarrow$ (rA)
$\mathrm{EA} \leftarrow \mathrm{b}+\mathrm{EXTZ}(\mathrm{UIMM} * 8)$
$\operatorname{MEM}(E A, 4) \leftarrow \mathrm{RS}_{0: 31}$
$\operatorname{MEM}(\mathrm{EA}+4,4) \leftarrow \mathrm{RS}_{32: 63}$

The contents of $\mathbf{r S}$ are stored as two words in storage addressed by EA.
Figure 5-139 shows how bytes are stored in memory as determined by the endian mode.

| GPR | a | b | C | d | e | f | g | h |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| Memory in big endian | a | b | C | d | e | f | g | h |
| Memory in little endian | d | c | b | a | h | g | f | e |

Figure 5-139. evstdw Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstdwx

| SPE | User |
| :--- | :--- |

Vector Store Double of Two Words Indexed
evstdwx rS,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrowb + (rB)
MEM(EA,4) \leftarrowRSS 0:31
MEM(EA+4,4) \leftarrow RS 32:63
```

The contents of $\mathbf{r S}$ are stored as two words in storage addressed by EA.
Figure 5-140 shows how bytes are stored in memory as determined by the endian mode.


| Byte address | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Memory in big endian | a | b | C | d | e | f | g | h |

Memory in little endian


Figure 5-140. evstdwx Results in Big- and Little-Endian Modes
Implementation note: If the EA is not double-word aligned, an alignment exception occurs.

## evstwhe

| SPE | User |
| :--- | :--- |

evstwhe
Vector Store Word of Two Half Words from Even
evstwhe rS,d(rA)


```
if (rA = 0) then b }\leftarrow
else b \leftarrow (rA)
EA \leftarrowb + EXTZ (UIMM*4)
MEM(EA,2) \leftarrow RS 0:15
MEM(EA+2,2) \leftarrow RS 32:47
```

The even half words from each element of $\mathbf{r S}$ are stored as two half words in storage addressed by EA.
Figure 5-141 shows how bytes are stored in memory as determined by the endian mode.


Figure 5-141. evstwhe Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwhex

Vector Store Word of Two Half Words from Even Indexed
evstwhex
rS,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
MEM(EA,2) \leftarrow RS 0:15
MEM (EA+2,2) \leftarrow RS 32:47
```

The even half words from each element of $\mathbf{r S}$ are stored as two half words in storage addressed by EA.
Figure 5-142 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |



Figure 5-142. evstwhex Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwho

SPE $\quad$ User

Vector Store Word of Two Half Words from Odd
evstwho rS,d(rA)


The odd half words from each element of $\mathbf{r S}$ are stored as two half words in storage addressed by EA, as shown in Figure 5-143.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | l

| Byte address | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| Memory in big endian | C | d | g | h |
| Memory in little endian | d | C | h | g |

Figure 5-143. evstwho Results in Big- and Little-Endian Modes

## evstwhox

Vector Store Word of Two Half Words from Odd Indexed
evstwhox rS,rA,rB

if $(r A=0)$ then $b \leftarrow 0$
else $\mathrm{b} \leftarrow(r \mathrm{~A})$
$\mathrm{EA} \leftarrow \mathrm{b}+(r \mathrm{~B})$
$\operatorname{MEM}(E A, 2) \leftarrow \mathrm{RS}_{16: 31}$
$\operatorname{MEM}(\mathrm{EA}+2,2) \leftarrow \mathrm{RS}_{48: 63}$
The odd half words from each element of $\mathbf{r S}$ are stored as two half words in storage addressed by EA.
Figure 5-144 shows how bytes are stored in memory as determined by the endian mode.


Figure 5-144. evstwhox Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwwe

| SPE | User |
| :--- | :--- |

## Vector Store Word of Word from Even

evstwwe rS,d(rA)


The even word of $\mathbf{r S}$ is stored in storage addressed by EA.
Figure 5-145 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |



Figure 5-145. evstwwe Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwwex

| SPE | User |
| :--- | :--- |

evstwwex
Vector Store Word of Word from Even Indexed
evstwwex
rS,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrowb + (rB)
MEM(EA,4)}\leftarrow~R\mp@subsup{S}{0:31}{
```

The even word of $\mathbf{r S}$ is stored in storage addressed by EA.
Figure 5-146 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

| Byte address | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| Memory in big endian | a | b | c | d |

Memory in little endian | d | c | b | a |
| :--- | :--- | :--- | :--- |

Figure 5-146. evstwwex Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwwo

SPE $\quad$ User

Vector Store Word of Word from Odd
evstwwo rS,d(rA)


The odd word of $\mathbf{r S}$ is stored in storage addressed by EA.
Figure 5-147 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |



Figure 5-147. evstwwo Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evstwwox

| SPE | User |
| :--- | :--- |

Vector Store Word of Word from Odd Indexed
evstwwox
rS,rA,rB


```
if (rA = 0) then b }\leftarrow
else b }\leftarrow(rA
EA \leftarrow b + (rB)
MEM(EA,4) \leftarrowrSS32:63
```

The odd word of $\mathbf{r S}$ is stored in storage addressed by EA.
Figure 5-148 shows how bytes are stored in memory as determined by the endian mode.

GPR | $a$ | $b$ | $c$ | $d$ | $e$ | $f$ | $g$ | $h$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |

| Byte address | 0 | 1 | 2 | 3 |
| :---: | :---: | :---: | :---: | :---: |
| Memory in big endian | e | f | g | h |

Memory in little endian | $h$ | $g$ | $f$ | $e$ |
| :--- | :--- | :--- | :--- |

Figure 5-148. evstwwox Results in Big- and Little-Endian Modes
Implementation note: If the EA is not word aligned, an alignment exception occurs.

## evsubfsmiaaw

| SPE | User |
| :--- | :--- |

Vector Subtract Signed, Modulo, Integer to Accumulator Word
evsubfsmiaaw rD,rA


$$
\begin{aligned}
& / / \text { high } \\
& r D_{0: 31} \leftarrow \operatorname{ACC}_{0: 31}-r A_{0: 31} \\
& / / \text { low } \\
& r D_{32: 63} \leftarrow A C C_{32: 63}-r A_{32: 63} \\
& / / \text { update accumulator } \\
& A C C_{0: 63} \leftarrow r D_{0: 63}
\end{aligned}
$$

Each word element in $\mathbf{r A}$ is subtracted from the corresponding element in the accumulator and the difference is placed into the corresponding rD word and into the accumulator, as shown in Figure 5-149.

Other registers altered: ACC


Figure 5-149. Vector Subtract Signed, Modulo, Integer to Accumulator Word (evsubfsmiaaw)

## evsubfssiaaw

| SPE | User |
| :--- | :--- |

## evsubfssiaaw

Vector Subtract Signed, Saturate, Integer to Accumulator Word
evsubfssiaaw rD,rA


Each signed integer word element in $\mathbf{r A}$ is sign-extended and subtracted from the corresponding sign-extended element in the accumulator, as shown in Figure 5-150, saturating if overflow occurs, and the results are placed in rD and the accumulator. Any overflow is recorded in the SPEFSCR overflow and summary overflow bits.

Other registers altered: SPEFSCR ACC


Figure 5-150. Vector Subtract Signed, Saturate, Integer to Accumulator Word (evsubfssiaaw)

## evsubfumiaaw

| SPE | User |
| :---: | :---: |

## evsubfumiaaw

Vector Subtract Unsigned, Modulo, Integer to Accumulator Word
evsubfumiaaw rD,rA


$$
\begin{aligned}
& / / \text { high } \\
& r D_{0: 31} \leftarrow \operatorname{ACC}_{0: 31}-r A_{0: 31} \\
& / / \text { low } \\
& r D_{32: 63} \leftarrow A C C_{32: 63}-r A_{32: 63} \\
& / / \text { update accumulator } \\
& A C C_{0: 63} \leftarrow r D_{0: 63}
\end{aligned}
$$

Each unsigned integer word element in $\mathbf{r A}$ is subtracted from the corresponding element in the accumulator and the results are placed in $\mathbf{r D}$ and into the accumulator, as shown in Figure 5-151.

Other registers altered: ACC


Figure 5-151. Vector Subtract Unsigned, Modulo, Integer to Accumulator Word (evsubfumiaaw)

## evsubfusiaaw

| SPE | User |
| :---: | :---: |

Vector Subtract Unsigned, Saturate, Integer to Accumulator Word
evsubfusiaaw rD,rA


Each unsigned integer word element in $\mathbf{r A}$ is zero-extended and subtracted from the corresponding zero-extended element in the accumulator, , as shown in Figure 5-152, saturating if underflow occurs, and the results are placed in $\mathbf{r D}$ and the accumulator. Any underflow is recorded in the SPEFSCR overflow and summary overflow bits.
Other registers altered: SPEFSCR ACC


Figure 5-152. Vector Subtract Unsigned, Saturate, Integer to Accumulator Word (evsubfusiaaw)

## evsubfw

Vector Subtract from Word
evsubfw rD,rA,rB


Each signed integer element of $\mathbf{r A}$ is subtracted from the corresponding element of $\mathbf{r B}$ and the results are placed into rD, as shown in Figure 5-153.


Figure 5-153. Vector Subtract from Word (evsubfw)

## evsubifw

| SPE | User |
| :--- | :--- |

## evsubifw

Vector Subtract Immediate from Word
evsubifw rD,UIMM,rB

$r D_{0: 31} \leftarrow r B_{0: 31}-\operatorname{EXTZ}(U I M M) / /$ Modulo difference
$r D_{32: 63} \leftarrow r B_{32: 63}-$ EXTZ(UIMM)// Modulo difference
UIMM is zero-extended and subtracted from both the high and low elements of $\mathbf{r B}$. Note that the same value is subtracted from both elements of the register, as shown in Figure 5-154. UIMM is 5 bits.


Figure 5-154. Vector Subtract Immediate from Word (evsubifw)

## evxor

| SPE | User |
| :--- | :--- |

Vector XOR
evxor rD,rA,rB

$r D_{0: 31} \leftarrow r A_{0: 31} \oplus r B_{0: 31} / /$ Bitwise XOR
$r D_{32: 63} \leftarrow r A_{32: 63} \oplus r B_{32: 63} / /$ Bitwise XOR
Each element of $\mathbf{r A}$ and $\mathbf{r B}$ is exclusive-ORed. The results are placed in $\mathbf{r D}$, as shown in Figure 5-155.


Figure 5-155. Vector XOR (evxor)

## Appendix A <br> Embedded Floating-Point Results Summary

Table A-1 through Table A-8 summarize the results of various types of embedded floating-point operations on various combinations of input operands. Flag settings are performed on appropriate element flags. For all the tables the following annotation and general rules apply:

-     * denotes that this status flag is set based on the results of the calculation.
- _Calc_ denotes that the result is updated with the results of the computation.
- max denotes the maximum normalized number with the sign set to the computation [sign(operand A) XOR sign(operand B)].
- amax denotes the maximum normalized number with the sign set to the sign of Operand A .
- bmax denotes the maximum normalized number with the sign set to the sign of Operand B.
- pmax denotes the maximum normalized positive number. The encoding for single-precision is: 0x7F7FFFFF. The encoding for double-precision is: 0x7FEFFFFF_FFFFFFFF.
- nmax denotes the maximum normalized negative number. The encoding for single-precision is: $0 x F F 7 F F F F F$. The encoding for double-precision is: 0xFFEFFFFF_FFFFFFFFF.
- pmin denotes the minimum normalized positive number. The encoding for single-precision is: $0 x 00800000$. The encoding for double-precision is: $0 x 00100000 \_00000000$.
- nmin denotes the minimum normalized negative number. The encoding for single-precision is: 0x80800000. The encoding for double-precision is: 0x80100000_00000000.
- Calculations that overflow or underflow saturate. Overflow for operations that have a floating-point result force the result to max. Underflow for operations that have a floating-point result force the result to zero. Overflow for operations that have a signed integer result force the result to 0x7FFFFFFF (positive) or 0x80000000 (negative). Overflow for operations that have an unsigned integer result force the result to 0xFFFFFFFF (positive) or 0x00000000 (negative).
- ${ }^{l}$ (superscript) denotes that the sign of the result is positive when the sign of Operand A and the sign of Operand B are different, for all rounding modes except round to minus infinity, where the sign of the result is then negative.
- ${ }^{2}$ (superscript) denotes that the sign of the result is positive when the sign of Operand A and the sign of Operand B are the same, for all rounding modes except round to minus infinity, where the sign of the result is then negative.
- ${ }^{3}$ (superscript) denotes that the sign for any multiply or divide is always the result of the operation [sign(Operand A) XOR sign(Operand B)].
- ${ }^{4}$ (superscript) denotes that if an overflow is detected, the result may be saturated.

Table A-1. Embedded Floating-Point Results Summary—Add, Sub, Mul, Div

| Operation | Operand A | Operand B | Result | FINV | FOVF | FUNF | FDBZ | FINX |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Add |  |  |  |  |  |  |  |  |  |
| Add | $\infty$ | $\infty$ | amax | 1 | 0 | 0 | 0 | 0 |  |
| Add | $\infty$ | NaN | amax | 1 | 0 | 0 | 0 | 0 |  |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Embedded Floating-Point Results Summary

Table A-1. Embedded Floating-Point Results Summary-Add, Sub, Mul, Div (continued)

| Operation | Operand A | Operand B | Result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Add | $\infty$ | denorm | amax | 1 | 0 | 0 | 0 | 0 |
| Add | $\infty$ | zero | amax | 1 | 0 | 0 | 0 | 0 |
| Add | $\infty$ | Norm | amax | 1 | 0 | 0 | 0 | 0 |
| Add | NaN | $\infty$ | amax | 1 | 0 | 0 | 0 | 0 |
| Add | NaN | NaN | amax | 1 | 0 | 0 | 0 | 0 |
| Add | NaN | denorm | amax | 1 | 0 | 0 | 0 | 0 |
| Add | NaN | zero | amax | 1 | 0 | 0 | 0 | 0 |
| Add | NaN | norm | amax | 1 | 0 | 0 | 0 | 0 |
| Add | denorm | $\infty$ | bmax | 1 | 0 | 0 | 0 | 0 |
| Add | denorm | NaN | bmax | 1 | 0 | 0 | 0 | 0 |
| Add | denorm | denorm | zero |  |  |  |  |  |
| 1 | 1 | 0 | 0 | 0 | 0 |  |  |  |
| Add | denorm | zero | zero |  |  |  |  |  |
| 1 | 1 | 0 | 0 | 0 | 0 |  |  |  |
| Add | denorm | norm | operand_b $^{4}$ | 1 | 0 | 0 | 0 | 0 |
| Add | zero | $\infty$ | bmax $^{2}$ | 1 | 0 | 0 | 0 | 0 |
| Add | zero | NaN | bmax $^{1}$ | 1 | 0 | 0 | 0 | 0 |
| Add | zero | denorm | zero |  |  |  |  |  |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table A-1. Embedded Floating-Point Results Summary-Add, Sub, Mul, Div (continued)

| Operation | Operand A | Operand B | Result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Sub | NaN | norm | amax | 1 | 0 | 0 | 0 | 0 |
| Sub | denorm | $\infty$ | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | denorm | NaN | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | denorm | denorm | zero ${ }^{2}$ | 1 | 0 | 0 | 0 | 0 |
| Sub | denorm | zero | zero ${ }^{2}$ | 1 | 0 | 0 | 0 | 0 |
| Sub | denorm | norm | -operand_b ${ }^{4}$ | 1 | 0 | 0 | 0 | 0 |
| Sub | zero | $\infty$ | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | zero | NaN | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | zero | denorm | zero ${ }^{2}$ | 1 | 0 | 0 | 0 | 0 |
| Sub | zero | zero | zero ${ }^{2}$ | 0 | 0 | 0 | 0 | 0 |
| Sub | zero | norm | -operand_b ${ }^{4}$ | 0 | 0 | 0 | 0 | 0 |
| Sub | norm | $\infty$ | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | norm | NaN | -bmax | 1 | 0 | 0 | 0 | 0 |
| Sub | norm | denorm | operand_a ${ }^{4}$ | 1 | 0 | 0 | 0 | 0 |
| Sub | norm | zero | operand_a ${ }^{4}$ | 0 | 0 | 0 | 0 | 0 |
| Sub | norm | norm | _Calc_ | 0 | * | * | 0 | * |
| Multiply ${ }^{3}$ |  |  |  |  |  |  |  |  |
| Mul | $\infty$ | $\infty$ | max | 1 | 0 | 0 | 0 | 0 |
| Mul | $\infty$ | NaN | max | 1 | 0 | 0 | 0 | 0 |
| Mul | $\infty$ | denorm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | $\infty$ | zero | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | $\infty$ | Norm | max | 1 | 0 | 0 | 0 | 0 |
| Mul | NaN | $\infty$ | max | 1 | 0 | 0 | 0 | 0 |
| Mul | NaN | NaN | max | 1 | 0 | 0 | 0 | 0 |
| Mul | NaN | denorm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | NaN | zero | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | NaN | norm | max | 1 | 0 | 0 | 0 | 0 |
| Mul | denorm | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | denorm | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | denorm | denorm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | denorm | zero | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | denorm | norm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | zero | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Embedded Floating-Point Results Summary

Table A-1. Embedded Floating-Point Results Summary-Add, Sub, Mul, Div (continued)

| Operation | Operand A | Operand B | Result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Mul | zero | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | zero | denorm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | zero | zero | zero | 0 | 0 | 0 | 0 | 0 |
| Mul | zero | norm | zero | 0 | 0 | 0 | 0 | 0 |
| Mul | norm | $\infty$ | $\max$ | 1 | 0 | 0 | 0 | 0 |
| Mul | norm | NaN | $\max$ | 1 | 0 | 0 | 0 | 0 |
| Mul | norm | denorm | zero | 1 | 0 | 0 | 0 | 0 |
| Mul | norm | zero | zero | 0 | 0 | 0 | 0 | 0 |
| Mul | norm | norm | _Calc_ | 0 | * | * | 0 | * |
| Divide ${ }^{3}$ |  |  |  |  |  |  |  |  |
| Div | $\infty$ | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Div | $\infty$ | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Div | $\infty$ | denorm | max | 1 | 0 | 0 | 0 | 0 |
| Div | $\infty$ | zero | max | 1 | 0 | 0 | 0 | 0 |
| Div | $\infty$ | Norm | max | 1 | 0 | 0 | 0 | 0 |
| Div | NaN | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Div | NaN | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Div | NaN | denorm | max | 1 | 0 | 0 | 0 | 0 |
| Div | NaN | zero | max | 1 | 0 | 0 | 0 | 0 |
| Div | NaN | norm | max | 1 | 0 | 0 | 0 | 0 |
| Div | denorm | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Div | denorm | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Div | denorm | denorm | max | 1 | 0 | 0 | 0 | 0 |
| Div | denorm | zero | max | 1 | 0 | 0 | 0 | 0 |
| Div | denorm | norm | zero | 1 | 0 | 0 | 0 | 0 |
| Div | zero | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Div | zero | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Div | zero | denorm | max | 1 | 0 | 0 | 0 | 0 |
| Div | zero | zero | max | 1 | 0 | 0 | 0 | 0 |
| Div | zero | norm | zero | 0 | 0 | 0 | 0 | 0 |
| Div | norm | $\infty$ | zero | 1 | 0 | 0 | 0 | 0 |
| Div | norm | NaN | zero | 1 | 0 | 0 | 0 | 0 |
| Div | norm | denorm | max | 1 | 0 | 0 | 0 | 0 |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table A-1. Embedded Floating-Point Results Summary—Add, Sub, Mul, Div (continued)

| Operation | Operand A | Operand B | Result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Div | norm | zero | max | 0 | 0 | 0 | 1 | 0 |
| Div | norm | norm | _Calc_ | 0 | $*$ | $*$ | 0 | $*$ |

Table A-2. Embedded Floating-Point Results Summary—Single Convert from Double

| Operand B | efscfd result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $+\infty$ | pmax | 1 | 0 | 0 | 0 | 0 |
| $-\infty$ | nmax | 1 | 0 | 0 | 0 | 0 |
| $+N a N$ | pmax | 1 | 0 | 0 | 0 | 0 |
| $-N a N$ | nmax | 1 | 0 | 0 | 0 | 0 |
| +denorm | +zero | 1 | 0 | 0 | 0 | 0 |
| -denorm | -zero | 1 | 0 | 0 | 0 | 0 |
| +zero | +zero | 0 | 0 | 0 | 0 | 0 |
| -zero | -zero | 0 | 0 | 0 | 0 | 0 |
| norm | -Calc_ | 0 | $*$ | $*$ | 0 | $*$ |

Table A-3. Embedded Floating-Point Results Summary—Double Convert from Single

| Operand B | efdcfs result | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $+\infty$ | pmax | 1 | 0 | 0 | 0 | 0 |
| $-\infty$ | nmax | 1 | 0 | 0 | 0 | 0 |
| $+N a N$ | pmax | 1 | 0 | 0 | 0 | 0 |
| $-N a N$ | nmax | 1 | 0 | 0 | 0 | 0 |
| +denorm | +zero | 1 | 0 | 0 | 0 | 0 |
| -denorm | -zero | 1 | 0 | 0 | 0 | 0 |
| +zero | +zero | 0 | 0 | 0 | 0 | 0 |
| -zero | -zero | 0 | 0 | 0 | 0 | 0 |
| norm | -Calc_ | 0 | 0 | 0 | 0 | 0 |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table A-4. Embedded Floating-Point Results Summary-Convert to Unsigned

| Operand B | Integer Result:ctui[d][z] | Fractional Result: ctuf | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $+\infty$ | 0xFFFF_FFFF <br> 0xFFFF_FFFF_FFFF_FFFF | 0x7FFF_FFFF | 1 | 0 | 0 | 0 | 0 |
| $-\infty$ | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| +NaN | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| -NaN | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| denorm | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| zero | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| + norm | _Calc_ | _Calc_ | $*$ | 0 | 0 | 0 | $*$ |
| - norm | _Calc_ | _Calc_ | $*$ | 0 | 0 | 0 | $*$ |

Table A-5. Embedded Floating-Point Results Summary-Convert to Signed

| Operand B | Integer Result <br> ctsi[d][z] | Fractional Result <br> ctsf | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $+\infty$ | 0x7FFF_FFFF <br> 0x7FFF_FFFF_FFFF_FFFF | 0x7FFF_FFFF | 1 | 0 | 0 | 0 | 0 |
| $-\infty$ | 0x8000_0000 <br> $0 \times 8000 \_0000 \_0000 \_0000$ | $0 \times 8000 \_0000$ | 1 | 0 | 0 | 0 | 0 |
| + NaN | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| $-N a N$ | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| denorm | 0 | 0 | 1 | 0 | 0 | 0 | 0 |
| zero | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| + norm | _Calc_ | _Calc_ | $*$ | 0 | 0 | 0 | $*$ |
| -norm | _Calc_ | _Calc_ | $*$ | 0 | 0 | 0 | $*$ |

Table A-6. Results Summary—Convert from Unsigned

| Operand B | Integer Source: cfui | Fractional Source: cfuf | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| zero | zero | zero | 0 | 0 | 0 | 0 | 0 |
| norm | _Calc_ | _Calc_ | 0 | 0 | 0 | 0 | $*$ |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table A-7. Embedded Floating-Point Results Summary—Convert from Signed

| Operand B | Integer Source: cfsi | Fractional Source: cfsf | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| zero | zero | zero | 0 | 0 | 0 | 0 | 0 |
| norm | _Calc_ | _Calc_ | 0 | 0 | 0 | 0 | $*$ |

Table A-8. Embedded Floating-Point Results Summary-*abs, *nabs, *neg

| Operand A | *abs | *nabs | *neg | FINV | FOVF | FUNF | FDBZ | FINX |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $+\infty$ | pmax ${ }^{+\infty}$ | $n \max \mid-\infty$ | $-\max \mid-\infty$ | 1 | 0 | 0 | 0 | 0 |
| $-\infty$ | pmax ${ }^{+\infty}$ | nmax ${ }^{-\infty}$ | -amax I $+\infty$ | 1 | 0 | 0 | 0 | 0 |
| +NaN | pmax ${ }^{\text {NaN }}$ | nmax $1-\mathrm{NaN}$ | -amax I-NaN | 1 | 0 | 0 | 0 | 0 |
| -NaN | pmax I NaN | nmax $1-\mathrm{NaN}$ | -amax I + NaN | 1 | 0 | 0 | 0 | 0 |
| +denorm | +zero I +denorm | -zero I-denorm | -zero I -denorm | 1 | 0 | 0 | 0 | 0 |
| -denorm | +zerol +denorm | -zero I -denorm | +zerol +denorm | 1 | 0 | 0 | 0 | 0 |
| +zero | +zero | -zero | -zero | 0 | 0 | 0 | 0 | 0 |
| -zero | +zero | -zero | +zero | 0 | 0 | 0 | 0 | 0 |
| +norm | +norm | -norm | -norm | 0 | 0 | 0 | 0 | 0 |
| -norm | +norm | -norm | +norm | 0 | 0 | 0 | 0 | 0 |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Appendix B SPE and Embedded Floating-Point Opcode Listings

This appendix lists SPE and embedded floating-point instructions as follows:

- Table B-1 lists opcodes alphabetically by mnemonic. Simplified mnemonics for SPE and embedded floating-point instructions are listed in this table with their standard instruction equivalents.
- Table B-2 lists opcodes in numerical order, showing both the decimal and the hexadecimal value for the primary opcodes.
- Table B-3 lists opcodes by form, showing the opcodes in binary.


## B. 1 Instructions (Binary) by Mnemonic

Table B-1 lists instructions by mnemonic.
Table B-1. Instructions (Binary) by Mnemonic


Table B-1. Instructions (Binary) by Mnemonic


| efdctuf | 0 | 001 | 00 | rD |  | /// | rB | 010 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | 011 |  | EFX efdctuf |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| efdctui | 0 | 001 | 00 | rD |  | /// | rB | 010 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | 010 | 0 | EFX efdctui |
| efdctuiz | 0 | 001 | 00 | rD |  | /// | rB | 010 | $1 \begin{array}{lll}1 & 1\end{array}$ | 100 | 0 | EFX efdctuiz |
| efddiv | 0 | 0001 | 00 | rD |  | rA | rB | 010 | 1110 | 100 | 1 | EFX efddiv |
| efdmul | 0 | 001 | 00 | rD |  | rA | rB | 010 | 1110 | 100 | 0 | EFX efdmul |
| efdnabs | 0 | 001 | 00 | rD |  | rA | /// | 010 | 1110 | 010 | 1 | EFX efdnabs |
| efdneg | 0 | 001 | 00 | rD |  | rA | /// | 010 | 1110 | 011 | 0 | EFX efdneg |
| efdsub | 0 | 001 | 00 | rD |  | rA | rB | 010 | 1110 | 000 | 1 | EFX efdsub |
| efdtsteq |  | 0001 | 00 | crfD | $1 /$ | rA | rB | 010 | 11111 | 111 | 0 | EFX efdtsteq |
| efdtstgt |  | 0001 | 00 | crfD | $1 /$ | rA | rB | 010 | 11111 | 110 | 0 | EFX efdtstgt |
| efdtstlt |  | 0001 | 00 | crfD | $1 /$ | rA | rB | 010 | 11111 | 110 | 1 | EFX efdtstlt |
| efsabs | 0 | 0001 | 00 | rD |  | rA | //I | 010 | 1100 | 010 | 0 | EFX efsabs |
| efsadd | 0 | 0001 | 00 | rD |  | rA | rB | 010 | 1100 | 000 | 0 | EFX efsadd |
| efscfd |  | 0001 | 00 | rD |  | 00000 | rB | 010 | 1100 | 111 | 1 | EFX efscfd |
| efscfsf |  | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 001 | 1 | EFX efscfsf |
| efscfsi | 0 | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 000 | 1 | EFX efscfsi |
| efscfuf | 0 | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 001 | 0 | EFX efscfuf |
| efscfui |  | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 000 | 0 | EFX efscfui |
| efscmpeq |  | 0001 | 00 | crfD | $1 /$ | rA | rB | 010 | 1100 | 111 | 0 | EFX efscmpeq |
| efscmpgt |  | 0001 | 00 | crfD | / / | rA | rB | 010 | 1100 | 110 | 0 | EFX efscmpgt |
| efscmplt |  | 0001 | 00 | crfD | / / | rA | rB | 010 | 1100 | 110 | 1 | EFX efscmplt |
| efsctsf |  | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 011 | 1 | EFX efsctsf |
| efsctsi |  | 0001 | 00 | rD |  | //I | rB | 010 | 1101 | 010 | 1 | EFX efsctsi |
| efsctsiz |  | 0001 | 00 | rD |  | /// | rB | 010 | 1101 | 101 | 0 | EFX efsctsiz |
| efsctuf |  | 0001 | 00 | rD |  | //] | rB | 010 | 1101 | 011 | 0 | EFX efsctuf |
| efsctui |  | 0001 | 00 | rD |  | /// | rB | 010 | 11001 | 010 | 0 | EFX efsctui |
| efsctuiz |  | 0001 | 00 | rD |  | //] | rB | 010 | 1101 | 100 | 0 | EFX efsctuiz |
| efsdiv |  | 0001 | 00 | rD |  | rA | rB | 010 | 1100 | 100 | 1 | EFX efsdiv |
| efsmul |  | 0001 | 00 | rD |  | rA | rB | 010 | 1100 | 100 | 0 | EFX efsmul |
| efsnabs |  | 0001 | 00 | rD |  | rA | //I | 010 | 1100 | 010 | 1 | EFX efsnabs |
| efsneg |  | 0001 | 00 | rD |  | rA | /// | 010 | 1100 | 011 | 0 | EFX efsneg |
| efssub |  | 0001 | 00 | rD |  | rA | rB | 010 | 1100 | 000 | 1 | EFX efssub |
| efststeq |  | 0001 | 00 | crfD | $1 /$ | rA | rB | 010 | 1101 | 111 | 0 | EFX efststeq |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-1. Instructions (Binary) by Mnemonic


Table B-1. Instructions (Binary) by Mnemonic



## B. 2 Instructions (Decimal and Hexadecimal) by Opcode

Table B-2 lists instructions by opcode.
Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


| evabs | 04 | rD |  | rA | //I | 010 | 0000 | 1000 | EVX evabs |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| evaddiw | 04 | rD |  | UIMM | rB | 010 | 0000 | 00010 | EVX evaddiw |
| evaddsmiaaw | 04 | rD |  | rA | /// | 100 | 1100 | 1001 | EVX evaddsmiaaw |
| evaddssiaaw | 04 | rD |  | rA | /// | 100 | 1100 | 00001 | EVX evaddssiaaw |
| evaddumiaaw | 04 | rD |  | rA | /// | 100 | 1100 | 1000 | EVX evaddumiaaw |
| evaddusiaaw | 04 | rD |  | rA | //I | 100 | 1100 | 0000 | EVX evaddusiaaw |
| evaddw | 04 | rD |  | rA | rB | 010 | 0000 | 0000 | EVX evaddw |
| evand | 04 | rD |  | rA | rB | 010 | 00001 | 00001 | EVX evand |
| evandc | 04 | rD |  | rA | rB | 010 | 00001 | 00010 | EVX evandc |
| evcmpeq <br> evcmpgts | 04 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 1 & 1\end{array}$ | 0100 | EVX evcmpeq |
|  | 04 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1\end{array}$ | $0 \quad 0001$ | EVX evcmpgts |
| evcmpgtu | 04 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1\end{array}$ | 0000 | EVX evcmpgtu |
| evcmplts | 04 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1\end{array}$ | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | EVX evcmplts |
| evcmpltu | 04 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | 00010 | EVX evcmpltu |
| evcntlsw | 04 | rD |  | rA | /// | 010 | 0000 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | EVX evcntlsw |
| evcntlzw | 04 | rD |  | rA | /// | 010 | 0000 | 1101 | EVX evcntlzw |
| evdivws | 04 | rD |  | rA | rB | 100 | 1100 | 01110 | EVX evdivws |
| evdivwu | 04 | rD |  | rA | rB | 100 | 1100 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX evdivwu |
| eveqv | 04 | rD |  | rA | rB | 010 | 00001 | 10001 | EVX eveqv |
| evextsb | 04 | rD |  | rA | /// | 010 | 0000 | 1010 | EVX evextsb |
| evextsh | 04 | rD |  | rA | /// | 010 | 0000 | $1 \begin{array}{llll}1 & 0 & 1 & 1\end{array}$ | EVX evextsh |
| evfsabs | 04 | rD |  | rA | /// | 010 | 1000 | 0100 | EVX evfsabs |
| evfsadd | 04 | rD |  | rA | rB | 010 | 1000 | 0000 | EVX evfsadd |
| evfscfsf | 04 | rD |  | /// | rB | 010 | 1001 | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | EVX evfscfsf |
| evfscfsi | 04 | rD |  | /// | rB | 010 | 10001 | 00001 | EVX evfscfsi |
| evfscfuf | 04 | rD |  | //] | rB | 010 | 1001 | 00010 | EVX evfscfuf |
| evfscfui | 04 | rD |  | //] | rB | 010 | 1001 | 0000 | EVX evfscfui |
| evfscmpeq | 04 | crfD | $1 /$ | rA | rB | 010 | 1000 | 1110 | EVX evfscmpeq |
| evfscmpgt | 04 | crfD | $1 /$ | rA | rB | 010 | 1000 | 1100 | EVX evfscmpgt |
| evfscmplt | 04 | crfD | $1 /$ | rA | rB | 010 | 1000 | 1101 | EVX evfscmplt |
| evfsctsf | 04 | rD |  | /// | rB | 010 | 1001 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX evfsctsf |
| evfsctsi | 04 | rD |  | /// | rB | 010 | 1001 | 01001 | EVX evfsctsi |
| evfsctsiz | 04 | rD |  | /// | rB | 010 | 1001 | 1010 | EVX evfsctsiz |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


| evfsctuf evfsctui | 04 | rD |  | /// | rB | 010 | 10001 | $\begin{array}{lllll}0 & 1 & 1 & 0\end{array}$ | EVX evfsctuf |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 04 | rD |  | /// | rB | 010 | 1001 | 0100 | EVX evfsctui |
| evfsctuiz | 04 | rD |  | /// | rB | 010 | 1001 | 1000 | EVX evfsctuiz |
| evfsdiv | 04 | rD |  | rA | rB | 010 | 1000 | 10001 | EVX evfsdiv |
| evfsmul | 04 | rD |  | rA | rB | 010 | 1000 | 1000 | EVX evfsmul |
| evfsnabs | 04 | rD |  | rA | /// | 010 | 1000 | $0 \begin{array}{llll}0 & 1 & 0 & 1\end{array}$ | EVX evfsnabs |
| evfsneg | 04 | rD |  | rA | /// | 010 | 1000 | $0 \begin{array}{llll}0 & 1 & 1\end{array}$ | EVX evfsneg |
| evfssub | 04 | rD |  | rA | rB | 010 | 1000 | 00001 | EVX evfssub |
| evfststeq | 04 | crfD | $1 /$ | rA | rB | 010 | 1001 | $1 \begin{array}{llll}1 & 1 & 0\end{array}$ | EVX evfststeq |
| evfststgt | 04 | crfD | $1 /$ | rA | rB | 010 | 1001 | 1100 | EVX evfststgt |
| evfststlt | 04 | crfD | $1 /$ | rA | rB | 010 | 1001 | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX evfststlt |
| efscfd | 04 | rD |  | 00000 | rB | 010 | 1100 | $\begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EFX efscfd |
| efdcfs | 04 | rD |  | 00000 | rB | 010 | 1110 | $\begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EFX efdcfs |
| evidd | 04 | rD |  | rA | UIMM ${ }^{1}$ | $0 \begin{array}{lll}0 & 1 & 1\end{array}$ | 0000 | 00001 | EVX evidd |
| eviddx | 04 | rD |  | rA | rB | 011 | 0000 | 0000 | EVX eviddx |
| evidh | 04 | rD |  | rA | UIMM ${ }^{1}$ | 011 | 0000 | $0 \begin{array}{llll}0 & 1 & 0 & 1\end{array}$ | EVX evidh |
| evidhx | 04 | rD |  | rA | rB | $0 \quad 11$ | 0000 | 01000 | EVX evidhx |
| evldw | 04 | rD |  | rA | UIMM ${ }^{1}$ | 011 | 0000 | $\begin{array}{lllll}0 & 0 & 1 & 1\end{array}$ | EVX evldw |
| evldwx | 04 | rD |  | rA | rB | 011 | 0000 | 00010 | EVX evldwx |
| evlhhesplat | 04 | rD |  | rA | UIMM ${ }^{2}$ | 0111 | 0000 | 10001 | EVX evlhhesplat |
| evlhhesplatx | 04 | rD |  | rA | rB | $0 \quad 11$ | 0000 | 1000 | EVX evlhhesplatx |
| evlhhossplat | 04 | rD |  | rA | UIMM ${ }^{2}$ | 0111 | 0000 | $\begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EVX evlhhossplat |
| evlhhossplatx | 04 | rD |  | rA | rB | 0111 | 0000 | 1110 | EVX evlhhossplatx |
| evlhhousplat | 04 | rD |  | rA | UIMM ${ }^{2}$ | 011 | 0000 | 1101 | EVX evlhhousplat |
| evIhhousplatx | 04 | rD |  | rA | rB | 011 | 0000 | 1100 | EVX evlhhousplatx |
| evlwhe | 04 | rD |  | rA | UIMM ${ }^{3}$ | $0 \quad 11$ | 00001 | 00001 | EVX evlwhe |
| evlwhex | 04 | rD |  | rA | rB | 0111 | 00001 | 0000 | EVX evlwhex |
| evlwhos | 04 | rD |  | rA | UIMM ${ }^{3}$ | 0111 | 00001 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX evlwhos |
| evlwhosx | 04 | rD |  | rA | rB | $0 \begin{array}{lll}0 & 1\end{array}$ | 00001 | $0 \begin{array}{llll}1 & 1 & 0\end{array}$ | EVX evlwhosx |
| evlwhou | 04 | rD |  | rA | UIMM $^{3}$ | 0111 | 00001 | 010011 | EVX evlwhou |
| evlwhoux | 04 | rD |  | rA | rB | $0 \quad 11$ | 00001 | 01000 | EVX evlwhoux |
| evlwhsplat | 04 | rD |  | rA | UIMM $^{3}$ | 011 | 00001 | $\begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX evlwhsplat |
| evlwhsplatx | 04 | rD |  | rA | rB | 011 | $0 \quad 0 \quad 0 \quad 1$ | 1100 | EVX evlwhsplatx |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-2. Instructions (Decimal and Hexadecimal) by Opcode



Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## SPE and Embedded Floating-Point Opcode Listings

## B. 3 Instructions by Form

Table B-3 lists instructions by form.
Table B-3. Instructions (Binary) by Form


## Table B-3. Instructions (Binary) by Form



| efscfui efscmpeq efscmpgt | $0 \quad 0001$ | 00 | rD |  | /// | rB | 010 | 11001 | 0000 | EFX efscfui <br> EFX efscmpeq <br> EFX efscmpgt |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $0 \quad 0 \quad 0 \quad 1$ | 00 | crfD | $1 /$ | rA | rB | 010 | 1100 | 1110 |  |
|  | 00001 | 00 | crfD | $1 / 1$ | rA | rB | 010 | 1100 | 1100 |  |
| efscmplt | 00001 | 00 | crfD | / / | rA | rB | 010 | 1100 | 1101 | EFX efscmplt |
| efsctsf | 000001 | 00 | rD |  | /// | rB | 010 | 11001 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EFX efsctsf |
| efsctsi | 00001 | 00 | rD |  | /// | rB | 010 | 11001 | $0 \begin{array}{llll}0 & 1 & 0 & 1\end{array}$ | EFX efsctsi |
| efsctsiz | 00001 | 00 | rD |  | /// | rB | 010 | 1101 | 1010 | EFX efsctsiz |
| efsctuf | 00001 | 00 | rD |  | //] | rB | 010 | 11001 | $\begin{array}{llll}0 & 1 & 1 & 0\end{array}$ | EFX efsctuf |
| efsctui | 000001 | 00 | rD |  | /// | rB | 010 | 11001 | 01000 | EFX efsctui |
| efsctuiz | 00001 | 00 | rD |  | /// | rB | 010 | 11001 | 1000 | EFX efsctuiz |
| efsdiv | 000001 | 00 | rD |  | rA | rB | 010 | 1100 | 1001 | EFX efsdiv |
| efsmul | 00001 | 00 | rD |  | rA | rB | 010 | 1100 | 1000 | EFX efsmul |
| efsnabs | 00001 | 00 | rD |  | rA | /// | 010 | 1100 | 010101 | EFX efsnabs |
| efsneg | 00001 | 00 | rD |  | rA | /// | 010 | 1100 | $0 \begin{array}{llll}0 & 1 & 1 & 0\end{array}$ | EFX efsneg |
| efssub | 00001 | 00 | rD |  | rA | rB | 010 | 1100 | 00001 | EFX efssub |
| efststeq | 00001 | 00 | crfD | $1 /$ | rA | rB | 010 | 1101 | $\begin{array}{llll}1 & 1 & 1 & 0\end{array}$ | EFX efststeq |
| efststgt | 00001 | 00 | crfD | $1 /$ | rA | rB | 010 | 1101 | 1100 | EFX efststgt |
| efststlt | 00001 | 00 | crfD | 11 | rA | rB | 010 | 1101 | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EFX efststit |
| brinc ${ }^{1}$ | 00001 | 00 | rD |  | rA | rB | 010 | 0000 | $\begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EVX brinc |
| evabs | 00001 | 00 | rD |  | rA | //I | 010 | 0000 | 1000 | EVX evabs |
| evaddiw | 00001 | 00 | rD |  | UIMM | rB | 010 | 0000 | 0010 | EVX evaddiw |
| evaddsmiaaw | 00001 | 00 | rD |  | rA | /// | 100 | 1100 | 10001 | EVX evaddsmiaaw |
| evaddssiaaw | 000001 | 00 | rD |  | rA | /// | 100 | 1100 | 0001 | EVX evaddssiaaw |
| evaddumiaaw | 00001 | 00 | rD |  | rA | /// | 100 | 1100 | 1000 | EVX evaddumiaaw |
| evaddusiaaw | 00001 | 00 | rD |  | rA | /// | 100 | 1100 | 0000 | EVX evaddusiaaw |
| evaddw | 00001 | 00 | rD |  | rA | rB | 010 | 0000 | 0000 | EVX evaddw |
| evand | 00001 | 00 | rD |  | rA | rB | 010 | 00001 | 00001 | EVX evand |
| evandc | 00001 | 00 | rD |  | rA | rB | 010 | 00001 | 00010 | EVX evandc |
| evcmpeq | 000001 | 00 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | 01000 | EVX evcmpeq |
| evcmpgts | 00001 | 00 | crfD | $1 /$ | rA | rB | 010 | 00011 | 00001 | EVX evcmpgts |
| evcmpgtu | 000001 | 00 | crfD | $1 /$ | rA | rB | 010 | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | 00000 | EVX evcmpgtu |
| evcmplts | 00001 | 00 | crfD | 11 | rA | rB | 010 | 00011 | 000111 | EVX evcmpls |
| evcmpltu | $0 \quad 001$ | 00 | crfD | $1 /$ | rA | rB | 010 | 00011 | 00010 | EVX evcmpltu |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-3. Instructions (Binary) by Form


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Table B-3. Instructions (Binary) by Form



Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Table B-3. Instructions (Binary) by Form

| Mnemonic | 0 | 123 | 45 | 8 | 1213 | 1617181920212223 |  | 24252627 | 28293031 | Form | Mnemonic |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| evmhesmi | 0 | 0001 | 00 | rD | rA | rB | 100 | 0000 | 10001 | EVX | evmhesmi |
| evmhesmia | 0 | 0001 | 00 | rD | rA | rB | 100 | 0010 | 10001 | EVX | evmhesmia |
| evmhesmiaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | 10001 | EVX | evmhesmiaaw |
| evmhesmianw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | 10001 | EVX | evmhesmianw |
| evmhessf | 0 | 0001 | 00 | rD | rA | rB | 100 | 0000 | 00011 | EVX | evmhessf |
| evmhessfa | 0 | 0001 | 00 | rD | rA | rB | 100 | 0010 | 00011 | EVX | evmhessfa |
| evmhessfaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | $0 \begin{array}{llll}0 & 0 & 1 & 1\end{array}$ | EVX | evmhessfaaw |
| evmhessfanw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | 00011 | EVX | evmhessfanw |
| evmhessiaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | 00001 | EVX | evmhessiaaw |
| evmhessianw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | 00001 | EVX | evmhessianw |
| evmheumi | 0 | 0001 | 00 | rD | rA | rB | 100 | 0000 | 1000 | EVX | evmheumi |
| evmheumia | 0 | 0001 | 00 | rD | rA | rB | 100 | 0010 | 1000 | EVX | evmheumia |
| evmheumiaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | 1000 | EVX | evmheumiaaw |
| evmheumianw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | 1000 | EVX | evmheumianw |
| evmheusiaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | 0000 | EVX | evmheusiaaw |
| evmheusianw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | 0000 | EVX | evmheusianw |
| evmhogsmfaa | 0 | 00001 | 00 | rD | rA | rB | 101 | 0010 | $1 \begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EVX | evmhogsmfaa |
| evmhogsmfan | 0 | 0001 | 00 | rD | rA | rB | 101 | 1010 | $1 \begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EVX | evmhogsmfan |
| evmhogsmiaa | 0 | 00001 | 00 | rD | rA | rB | 101 | $0 \quad 0 \quad 10$ | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX | evmhogsmiaa |
| evmhogsmian | 0 | 0001 | 00 | rD | rA | rB | 101 | 1010 | 11 1  | EVX | evmhogsmian |
| evmhogumiaa | 0 | 0001 | 00 | rD | rA | rB | 101 | 00010 | 1100 | EVX | evmhogumiaa |
| evmhogumian | 0 | 0001 | 00 | rD | rA | rB | 101 | 1010 | 1100 | EVX | evmhogumian |
| evmhosmf | 0 | 0001 | 00 | rD | rA | rB | 100 | 0000 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | EVX | evmhosmf |
| evmhosmfa | 0 | 0001 | 00 | rD | rA | rB | 100 | 0010 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | EVX | evmhosmfa |
| evmhosmfaaw | 0 | 0001 | 00 | rD | rA | rB | 101 | 0000 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | EVX | evmhosmfaaw |
| evmhosmfanw | 0 | 0001 | 00 | rD | rA | rB | 101 | 1000 | $\begin{array}{llll}1 & 1 & 1 & 1\end{array}$ | EVX | evmhosmfanw |
| evmhosmi | 0 | 0001 | 00 | rD | rA | rB | 100 | 0000 | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX | evmhosmi |
| evmhosmia | 0 | 0001 | 00 | rD | rA | rB | 100 | 0010 | $1 \begin{array}{llll}1 & 1 & 1\end{array}$ | EVX | evmhosmia |
| evmhosmiaaw |  | 0001 | 00 | rD | rA | rB | 101 | 0000 | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX | evmhosmiaaw |
| evmhosmianw |  | 0001 | 00 | rD | rA | rB | 101 | 1000 | $1 \begin{array}{llll}1 & 1 & 0 & 1\end{array}$ | EVX | evmhosmianw |
| evmhossf |  | 0001 | 00 | rD | rA | rB | 100 | 0000 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX | evmhossf |
| evmhossfa |  | 0001 | 00 | rD | rA | rB | 100 | 0010 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX | evmhossfa |
| evmhossfaaw |  | 0001 | 00 | rD | rA | rB | 101 | 0000 | $\begin{array}{lllll}0 & 1 & 1 & 1\end{array}$ | EVX | evmhossfaaw |

Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

## Table B-3. Instructions (Binary) by Form



Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-3. Instructions (Binary) by Form


Signal Processing Engine (SPE) Programming Environments Manual, Rev. 0

Table B-3. Instructions (Binary) by Form


SPE and Embedded Floating-Point Opcode Listings

