

## **Introduction**

This reference manual gives an overview of Book E, a version of the PowerPC architecture intended for embedded processors. To ensure application level compatibility with the PowerPC architecture developed by Apple, IBM, and Freescale, Book E incorporates the user level resources defined in the user instruction set architecture (UISA), Book I, of the AIM architectural definition.

This reference manual is a reference document for the SPC560x and RPC560Bx devices which are based on the PowerPC architecture.

# Contents

<b>1</b>	<b>Preface</b> .....	<b>25</b>
1.1	About this book .....	25
1.1.1	Audience .....	25
1.1.2	Organization .....	26
1.2	Suggested reading .....	27
1.2.1	General information .....	27
1.2.2	Related documentation .....	27
1.2.3	Conventions .....	28
1.3	Acronyms and abbreviations .....	29
1.4	Terminology conventions .....	30
1.5	Part I: Book E and Book E implementation standards .....	31
1.6	Part II: EIS-defined extensions to the Book E architecture .....	32
<b>2</b>	<b>Overview</b> .....	<b>33</b>
2.1	Overview Book E and the Book E implementation standards (EIS) .....	33
2.1.1	Auxiliary processing units (APUs) .....	34
2.2	Instruction set .....	34
2.3	Register set .....	35
2.4	Interrupts and exception handling .....	36
2.4.1	Exception handling .....	36
2.4.2	Interrupt classes .....	36
2.4.3	Interrupt categories .....	36
2.4.4	Interrupt registers .....	37
2.5	Memory management .....	40
2.5.1	Address translation .....	40
2.5.2	MMU assist registers (MAS1–MAS7) .....	41
2.5.3	Process ID registers (PID0–PID2) .....	42
2.5.4	TLB coherency .....	42
2.5.5	Atomic update memory references .....	42
2.5.6	Memory access ordering .....	42
2.5.7	Cache control instructions .....	43
2.5.8	Programmable page characteristics .....	43
2.6	Performance monitoring .....	43

2.6.1	Global control register	43
2.6.2	Performance monitor counter registers	43
2.6.3	Local control registers	44
2.7	Legacy support of PowerPC architecture	44
2.7.1	Instruction set compatibility	44
2.7.2	Memory subsystem	45
2.7.3	Interrupt handling	45
2.7.4	Memory management	45
2.7.5	Requirements for system reset generation	45
2.7.6	Little-endian mode	45
<b>3</b>	<b>Register model</b>	<b>46</b>
3.1	Overview	46
3.2	Register model for 32-bit Book E implementations	47
3.2.1	Special-purpose registers (SPRs)	50
3.3	Registers for integer operations	55
3.3.1	General purpose registers (GPRs)	56
3.3.2	Integer exception register (XER)	57
3.4	Registers for floating-point operations	59
3.4.1	Floating-point registers (FPRs)	59
3.4.2	Floating-point status and control register (FPSCR)	59
3.5	Registers for branch operations	62
3.5.1	Condition register (CR)	62
3.5.2	Link register (LR)	68
3.5.3	Count register (CTR)	68
3.6	Processor control registers	69
3.6.1	Machine state register (MSR)	69
3.7	Hardware implementation-dependent registers	72
3.7.1	Hardware implementation dependent register 0 (HID0)	72
3.7.2	Hardware implementation dependent register 1 (HID1)	75
3.7.3	Processor ID register (PIR)	76
3.7.4	Processor version register (PVR)	76
3.7.5	System version register (SVR)	77
3.8	Timer registers	78
3.8.1	Timer control register (TCR)	79
3.8.2	Timer status register (TSR)	80

3.8.3	Time base (TBU and TBL) . . . . .	81
3.8.4	Decrementer register . . . . .	82
3.8.5	Decrementer auto-reload register (DECAR) . . . . .	83
3.9	Interrupt registers . . . . .	83
3.9.1	Interrupt registers defined by book E . . . . .	83
3.10	Software use sprs (SPRG0–SPRG7 and USPRG0) . . . . .	92
3.11	L1 cache registers . . . . .	92
3.11.1	L1 cache control and status register 0 (L1CSR0) . . . . .	92
3.11.2	L1 cache control and status register 1 (L1CSR1) . . . . .	94
3.11.3	L1 cache configuration register 0 (L1CFG0) . . . . .	96
3.11.4	L1 cache configuration register 1 (L1CFG1) . . . . .	97
3.11.5	L1 flush and invalidate control register 0 (L1FINV0) . . . . .	98
3.12	MMU registers . . . . .	99
3.12.1	Process ID registers (PID0–PIDn) . . . . .	99
3.12.2	MMU control and status register 0 (MMUCSR0) . . . . .	100
3.12.3	MMU configuration register (MMUCFG) . . . . .	101
3.12.4	TLB configuration registers (TLBnCFG) . . . . .	102
3.12.5	MMU assist registers (MAS0–MAS7) . . . . .	103
3.13	Debug registers . . . . .	109
3.13.1	Debug control registers (DBCR0–DBCR3) . . . . .	110
3.13.2	Debug status register (DBSR) . . . . .	117
3.13.3	Instruction address compare registers (IAC1–IAC4) . . . . .	118
3.13.4	Data address compare registers (DAC1–DAC2) . . . . .	119
3.13.5	Data value compare registers (DVC1 and DVC2) . . . . .	119
3.14	SPE and SPFP APU registers . . . . .	119
3.14.1	Signal processing, embedded floating-point status, control register (SPEFSCR) . . . . .	120
3.14.2	Accumulator (ACC) . . . . .	123
3.15	Alternate time base registers (ATBL and ATBU) . . . . .	123
3.16	Performance monitor registers (PMRs) . . . . .	124
3.16.1	Global control register 0 (PMGC0) . . . . .	126
3.16.2	User global control register 0 (UPMGC0) . . . . .	127
3.16.3	Local control A registers (PMLCa0–PMLCa3) . . . . .	127
3.16.4	User local control A registers (UPMLCa0–UPMLCa3) . . . . .	128
3.16.5	Local control B registers (PMLCb0–PMLCb3) . . . . .	128
3.16.6	User local control B registers (UPMLCb0–UPMLCb3) . . . . .	130

3.16.7	Performance monitor counter registers (PMC0–PMC3)	130
3.16.8	User performance monitor counter registers (UPMC0–UPMC3)	130
3.17	Device control registers (DCRs)	131
3.18	Book E SPR model	131
3.18.1	Invalid SPR references	131
3.18.2	Synchronization requirements for SPRs	131
3.18.3	Reserved SPRs	132
3.18.4	Allocated SPRs	132
<b>4</b>	<b>Instruction model</b>	<b>134</b>
4.1	Operand conventions	134
4.1.1	Data organization in memory and data transfers	134
4.1.2	Alignment and misaligned accesses	134
4.2	Instruction set summary	135
4.2.1	Classes of instructions	136
4.2.2	Instruction forms	140
4.2.3	Addressing modes	140
4.3	Instruction set overview	148
4.3.1	Book E user-level instructions	148
4.3.2	Supervisor level instructions	185
4.3.3	Recommended simplified mnemonics	188
4.3.4	Book E instructions with implementation-specific features	188
4.3.5	EIS instructions	189
4.3.6	Context synchronization	189
4.4	Instruction fetching	189
4.5	Memory synchronization	190
4.6	EIS-specific instructions	190
4.6.1	SPE and embedded floating-point APUs	190
4.6.2	Integer select ( <b>isel</b> ) APU	201
4.6.3	Performance monitor APU	201
4.6.4	Cache locking APU	203
4.6.5	Machine check APU	205
4.6.6	VLE extension	205
4.7	Instruction listing	236
<b>5</b>	<b>Interrupts and exceptions</b>	<b>250</b>

5.1	Overview .....	250
5.2	EIs interrupt definitions .....	251
5.2.1	Recoverability from interrupts .....	252
5.3	Interrupt registers .....	252
5.4	Exceptions .....	255
5.5	Interrupt classes .....	256
5.5.1	Requirements for system reset generation .....	257
5.6	Interrupt processing .....	258
5.7	Interrupt definitions .....	259
5.7.1	Critical input interrupt .....	261
5.7.2	Machine check interrupt .....	262
5.7.3	Data storage interrupt .....	263
5.7.4	Instruction storage interrupt .....	265
5.7.5	External input interrupt .....	266
5.7.6	Alignment interrupt .....	267
5.7.7	Program interrupt .....	268
5.7.8	Floating-point unavailable interrupt .....	270
5.7.9	System call interrupt .....	271
5.7.10	Auxiliary processor unavailable interrupt .....	271
5.7.11	Decrementer Interrupt .....	271
5.7.12	Fixed-interval timer interrupt .....	272
5.7.13	Watchdog timer interrupt .....	273
5.7.14	Data tlb error interrupt .....	273
5.7.15	Instruction tlb error interrupt .....	274
5.7.16	Debug interrupt .....	275
5.7.17	EIS-defined interrupts .....	275
5.8	Performance monitor interrupt .....	277
5.9	Partially executed instructions .....	277
5.10	Interrupt ordering and masking .....	279
5.10.1	Guidelines for system software .....	279
5.10.2	Interrupt order .....	281
5.11	Exception priorities .....	281
<b>6</b>	<b>Storage architecture .....</b>	<b>286</b>
6.1	<b>Overview .....</b>	<b>286</b>
6.2	Memory and cache coherency .....	286

6.2.1	Memory/Cache access attributes	287
6.2.2	Shared memory	294
6.3	Cache model	300
6.3.1	Cache programming model	301
6.3.2	Primary (L1) cache model	305
6.4	Storage model	306
6.4.1	Storage programming model	306
6.4.2	The storage architecture	308
6.4.3	Virtual address (VA)	309
6.4.4	Address spaces	309
6.4.5	Process ID	311
6.4.6	Address translation	312
6.4.7	Address translation and the ST EIS	314
6.4.8	Permission attributes	319
6.4.9	Translation lookaside buffer (TLB) arrays	321
6.4.10	TLB management	323
6.4.11	MAS registers and exception handling	329
<b>7</b>	<b>Instruction set</b>	<b>334</b>
7.1	Notation	334
7.2	Instruction fields	335
7.3	Description of instruction operations	336
7.3.1	SPE APU saturation and bit-reverse models	340
7.3.2	Embedded floating-point conversion models	341
7.3.3	Integer saturation models	350
7.3.4	Embedded floating-point results	351
7.4	Instruction set	351
<b>8</b>	<b>Auxiliary processing units (APUs)</b>	<b>667</b>
8.1	Integer select APU	667
8.1.1	Integer select APU programming model	667
8.1.2	Using <b>isel</b> to Improve conditional branch performance	668
8.2	Performance monitor APU	668
8.2.1	Performance monitor APU programming model	668
8.3	Signal processing engine APU (SPE APU)	670
8.3.1	Overview	670

8.3.2	Nomenclature and conventions	671
8.3.3	Programming model	671
8.3.4	Instruction definitions	676
8.4	Embedded vector and scalar single-precision floating-point APUs (SPFP APUs)	676
8.4.1	Nomenclature and conventions	676
8.4.2	Embedded floating-point APUs programming model	676
8.4.3	Embedded floating-point APU operations	683
8.4.4	Implementation options summary	686
8.5	Machine check APU	687
8.5.1	Machine check APU programming model	687
8.6	Debug APU	688
8.6.1	Debug APU programming model	688
8.6.2	Debug APU register model	688
8.6.3	Debug APU instruction model	689
8.7	Alternate time base	690
8.7.1	Programming model	690
<b>9</b>	<b>Storage-related APUs</b>	<b>692</b>
9.1	Cache line locking APU	692
9.1.1	Programming model	692
9.2	Direct cache flush APU	694
9.2.1	Overview	694
9.2.2	Programming model	694
9.3	Cache way partitioning APU	695
9.3.1	Programming model	695
9.3.2	Interaction with the cache locking APU	695
<b>10</b>	<b>VLE introduction</b>	<b>697</b>
10.1	Compatibility with PowerPC Book E	697
10.2	Instruction mnemonics and operands	698
<b>11</b>	<b>VLE storage addressing</b>	<b>699</b>
11.1	Data memory addressing modes	699
11.2	Instruction memory addressing modes	699



<b>12</b>	<b>VLE compatibility with the EIS</b> .....	<b>701</b>
12.1	Overview .....	701
12.2	VLE extension processor and storage control extensions .....	701
12.2.1	EIS instruction extensions .....	701
12.2.2	Book E instruction extensions .....	702
12.2.3	EIS MMU extensions .....	702
12.2.4	EIS debug APU extensions .....	704
<b>13</b>	<b>VLE instruction classes</b> .....	<b>705</b>
13.1	Processor control instructions .....	705
13.1.1	System linkage instructions .....	705
13.1.2	Processor control register manipulation instructions .....	705
13.1.3	Instruction synchronization instruction .....	706
13.2	Branch operation instructions .....	706
13.2.1	Registers for branch operations .....	706
13.2.2	Branch instructions .....	709
13.3	Condition register instructions .....	710
13.4	Integer instructions .....	711
13.4.1	Integer load instructions .....	711
13.4.2	Integer store instructions .....	712
13.4.3	Integer arithmetic instructions .....	713
13.4.4	Integer logical and move instructions .....	715
13.4.5	Integer compare and bit test instructions .....	717
13.4.6	Integer select instruction .....	718
13.4.7	Integer trap instructions .....	719
13.4.8	Integer rotate and shift instructions .....	719
13.5	Storage control instructions .....	721
13.5.1	Storage synchronization instructions .....	721
13.5.2	Cache management instructions .....	721
13.5.3	TLB management instructions .....	722
13.5.4	Instruction alignment and byte ordering .....	722
13.6	Instruction listings .....	722
<b>14</b>	<b>VLE instruction set</b> .....	<b>736</b>
14.1	Book E– and EIS-defined instructions .....	736
14.2	Immediate field and displacement field encodings .....	740

<b>15</b>	<b>VLE instruction index</b> .....	<b>778</b>
15.1	Instruction index sorted by opcode .....	778
15.2	Instruction index sorted by mnemonic .....	798
15.3	Instruction index sorted by opcode .....	817
15.4	Instruction index sorted by mnemonic .....	834
<b>Appendix A</b>	<b>Instruction set listings</b> .....	<b>850</b>
A.1	Instructions sorted by mnemonic (decimal and hexadecimal) .....	850
A.2	Instructions sorted by primary opcodes (decimal and hexadecimal) . . .	884
A.3	Instructions sorted by mnemonic (binary) .....	903
A.4	Instructions sorted by opcode (binary) .....	930
A.5	Instruction set legend .....	951
<b>Appendix B</b>	<b>Simplified mnemonics for PowerPC instructions</b> .....	<b>963</b>
B.1	Overview .....	963
B.2	Subtract simplified mnemonics .....	963
B.2.1	Subtract immediate .....	963
B.2.2	Subtract .....	964
B.3	Rotate and shift simplified mnemonics .....	964
B.3.1	Operations on words .....	964
B.4	Branch instruction simplified mnemonics .....	965
B.4.1	Key facts about simplified branch mnemonics .....	967
B.4.2	Eliminating the BO operand .....	967
B.4.3	Incorporating the BO branch prediction .....	969
B.4.4	The BI operand—CR bit and field representations .....	969
B.4.5	Simplified mnemonics that incorporate the BO operand .....	973
B.4.6	Simplified mnemonics that incorporate CR conditions (eliminates BO and replaces BI with <b>crS</b> ) .....	976
B.5	Compare word simplified mnemonics .....	980
B.6	Condition register logical simplified mnemonics .....	981
B.7	Trap instructions simplified mnemonics .....	982
B.8	Simplified mnemonics for accessing SPRs .....	983
B.9	Recommended simplified mnemonics .....	984
B.9.1	No-op ( <b>nop</b> ) .....	984
B.9.2	Load immediate ( <b>li</b> ) .....	984

B.9.3	Load address ( <b>la</b> ) . . . . .	985
B.9.4	Move register ( <b>mr</b> ) . . . . .	985
B.9.5	Complement register ( <b>not</b> ) . . . . .	985
B.9.6	Move to condition register ( <b>mtcr</b> ) . . . . .	985
B.10	EIS-specific simplified mnemonics . . . . .	985
B.10.1	Integer select ( <b>isel</b> ) . . . . .	986
B.10.2	SPE mnemonics . . . . .	986
B.11	Comprehensive list of simplified mnemonics . . . . .	986
<b>Appendix C Programming examples . . . . .</b>		<b>995</b>
C.1	Synchronization. . . . .	995
C.1.1	Synchronization primitives . . . . .	995
C.1.2	Lock acquisition and release . . . . .	997
C.1.3	List insertion . . . . .	998
C.1.4	Synchronization notes . . . . .	999
C.2	Multiple-precision shifts. . . . .	1000
C.3	Floating point conversions. . . . .	1001
C.3.1	Conversion from floating-point number to signed integer word . . . . .	1002
C.3.2	Conversion from floating-point number to unsigned integer word . . . . .	1003
C.4	Floating point selection . . . . .	1003
C.4.1	Notes . . . . .	1004
<b>Appendix D Guidelines for 32-bit book E . . . . .</b>		<b>1006</b>
D.1	Registers on 32-bit book E implementations. . . . .	1006
D.2	Addressing on 32-bit book E implementations . . . . .	1006
D.3	TLB fields on 32-bit book E implementations . . . . .	1006
D.4	32-bit book E software guidelines . . . . .	1007
D.4.1	32-bit instruction selection . . . . .	1007
D.4.2	32-bit addressing . . . . .	1007
<b>Appendix E Embedded floating-point results . . . . .</b>		<b>1008</b>
E.1	Notation conventions and general rules . . . . .	1008
E.2	Add, subtract, multiply, and divide results . . . . .	1009
E.3	Double- to single-precision conversion . . . . .	1012
E.4	Single- to double-precision conversion . . . . .	1013
E.5	Conversion to unsigned. . . . .	1013



---

E.6	Conversion to signed. . . . .	1014
E.7	Conversion from unsigned . . . . .	1014
E.8	Conversion from signed . . . . .	1014
E.9	*abs, *nabs, and *neg operations . . . . .	1015
<b>Glossary . . . . .</b>		<b>1016</b>
<b>Revision history . . . . .</b>		<b>1024</b>

## List of tables

Table 1.	Conventions . . . . .	28
Table 2.	Acronyms and abbreviated terms . . . . .	29
Table 3.	Terminology conventions . . . . .	30
Table 4.	Instruction field conventions . . . . .	31
Table 5.	Interrupt registers . . . . .	37
Table 6.	Interrupt vector registers and exception conditions . . . . .	39
Table 7.	Book E special purpose registers (by SPR abbreviation) . . . . .	50
Table 8.	EIS-defined SPRs (by SPR abbreviation) . . . . .	53
Table 9.	XER field descriptions . . . . .	58
Table 10.	FPSCR field descriptions . . . . .	60
Table 11.	Floating-point result flags . . . . .	62
Table 12.	BI operand settings for CR fields . . . . .	63
Table 13.	CR0 bit descriptions . . . . .	64
Table 14.	CR setting for floating-point instructions . . . . .	65
Table 15.	CR setting for compare instructions . . . . .	65
Table 16.	CR0 encodings . . . . .	67
Table 17.	Condition register setting for compare instructions . . . . .	67
Table 18.	Branch to link register instruction comparison . . . . .	68
Table 19.	Branch to count register instruction comparison . . . . .	69
Table 20.	MSR field descriptions . . . . .	70
Table 21.	Floating-point exception bits—MSR[FE0,FE1] . . . . .	72
Table 22.	HID0 field descriptions . . . . .	73
Table 23.	PVR field descriptions . . . . .	76
Table 24.	TCR field descriptions . . . . .	79
Table 25.	TSR field descriptions . . . . .	80
Table 26.	IVOR assignments . . . . .	86
Table 27.	Exception syndrome register (ESR) definition . . . . .	87
Table 28.	MCSR field descriptions . . . . .	91
Table 29.	L1CSR0 field descriptions . . . . .	93
Table 30.	L1CSR1 field descriptions . . . . .	95
Table 31.	L1CFG0 field descriptions . . . . .	96
Table 32.	L1CFG1 field descriptions . . . . .	98
Table 33.	L1FINV0 fields—L1 direct cache flush . . . . .	99
Table 34.	MMUCSR0 field descriptions . . . . .	100
Table 35.	MMUCFG field descriptions . . . . .	101
Table 36.	TLBnCFG field descriptions . . . . .	102
Table 37.	MAS0 field descriptions . . . . .	103
Table 38.	MAS1 field descriptions—descriptor context and configuration control . . . . .	104
Table 39.	MAS2 field descriptions—EPN and page attributes . . . . .	105
Table 40.	MAS3 field descriptions—RPN and access control . . . . .	106
Table 41.	MAS4 field descriptions—hardware replacement assist configuration . . . . .	107
Table 42.	MAS5 field descriptions—extended search pIDs . . . . .	108
Table 43.	MAS 6 field descriptions—search pids and search AS . . . . .	108
Table 44.	MAS 7 field descriptions—high order RPN . . . . .	109
Table 45.	DBCRO field descriptions . . . . .	110
Table 46.	DBCRI field descriptions . . . . .	112
Table 47.	DBCRI field descriptions . . . . .	115
Table 48.	DBSR field descriptions . . . . .	117

Table 49.	SPEFSCR field descriptions . . . . .	120
Table 50.	ATBL field descriptions . . . . .	124
Table 51.	ATBU field descriptions . . . . .	124
Table 52.	Performance monitor registers—supervisor level . . . . .	124
Table 53.	Performance monitor registers—user level (read-only) . . . . .	125
Table 54.	PMGC0 field descriptions . . . . .	126
Table 55.	PMLCa0–PMLCa3 field descriptions . . . . .	128
Table 56.	PMLCb0 –PMLCb3 field descriptions . . . . .	129
Table 57.	PMC0–PMC3 field descriptions . . . . .	130
Table 58.	System response to an invalid spr reference . . . . .	131
Table 59.	Synchronization requirements for sprs . . . . .	132
Table 60.	Allocated SPRs defined by the EIS . . . . .	132
Table 61.	Address characteristics of aligned operands . . . . .	135
Table 62.	Allocated instructions . . . . .	137
Table 63.	Preserved instructions . . . . .	138
Table 64.	Synchronization requirements . . . . .	144
Table 65.	Integer arithmetic instructions . . . . .	148
Table 66.	Integer 32-Bit compare instructions (L = 0) . . . . .	149
Table 67.	Integer logical instructions . . . . .	150
Table 68.	Integer rotate instructions . . . . .	151
Table 69.	Integer shift instructions . . . . .	151
Table 70.	Floating-point load instruction set . . . . .	153
Table 71.	Floating-point store instructions . . . . .	154
Table 72.	Floating-point move instructions . . . . .	155
Table 73.	Floating-point elementary arithmetic instructions . . . . .	155
Table 74.	Floating-point multiply-add instructions . . . . .	156
Table 75.	Floating-point rounding and conversion instructions . . . . .	156
Table 76.	CR field settings . . . . .	157
Table 77.	Floating-point compare and select instructions . . . . .	157
Table 78.	Floating-point status and control register instructions . . . . .	158
Table 79.	Integer load instructions . . . . .	161
Table 80.	Integer store instructions . . . . .	162
Table 81.	Integer load and store with byte-reverse instructions . . . . .	162
Table 82.	Integer load and store multiple instructions . . . . .	163
Table 83.	Integer load and store string instructions . . . . .	163
Table 84.	Floating-point load instructions . . . . .	163
Table 85.	Floating-point store instructions . . . . .	164
Table 86.	Store floating-point single behavior . . . . .	164
Table 87.	Store floating-point double behavior . . . . .	165
Table 88.	BO bit descriptions . . . . .	169
Table 89.	BO operand encodings . . . . .	170
Table 90.	Branch instructions . . . . .	171
Table 91.	Condition register logical instructions . . . . .	171
Table 92.	Trap instructions . . . . .	172
Table 93.	System linkage instruction . . . . .	172
Table 94.	Move to/from condition register instructions . . . . .	172
Table 95.	Move to/from special-purpose register instructions . . . . .	173
Table 96.	Book E special-purpose registers (by SPR abbreviation) . . . . .	173
Table 97.	Implementation-specific SPRs (by SPR abbreviation) . . . . .	176
Table 98.	Memory synchronization instructions . . . . .	178
Table 99.	User-level cache instructions . . . . .	184
Table 100.	System linkage instructions—supervisor-level . . . . .	185

Table 101.	Move to/from machine state register instructions . . . . .	186
Table 102.	Supervisor-Level cache management instruction . . . . .	186
Table 103.	TLB management instructions . . . . .	187
Table 104.	Implementation-specific instructions summary . . . . .	188
Table 105.	EIS-defined instructions (except SPE and SPFP instructions) . . . . .	189
Table 106.	SPE APU vector multiply instruction mnemonic structure . . . . .	192
Table 107.	Mnemonic extensions for multiply-accumulate instructions . . . . .	192
Table 108.	SPE APU vector instructions . . . . .	193
Table 109.	Vector and scalar floating-point APU instructions . . . . .	200
Table 110.	Integer select APU instruction . . . . .	201
Table 111.	Performance monitor APU instructions . . . . .	201
Table 112.	Performance monitor registers—supervisor level . . . . .	202
Table 113.	Performance monitor registers—user level (read-only) . . . . .	203
Table 114.	Cache locking APU instructions . . . . .	204
Table 115.	Machine check APU instruction . . . . .	205
Table 116.	System linkage instruction set index . . . . .	205
Table 117.	System register manipulation instruction set index . . . . .	206
Table 118.	Instruction Synchronization Instruction Set Index . . . . .	206
Table 119.	VLE extension BO32 encodings . . . . .	207
Table 120.	VLE extension BO16 encodings . . . . .	207
Table 121.	Branch instruction set index . . . . .	208
Table 122.	Condition register instruction set index . . . . .	208
Table 123.	Basic integer load instruction set index . . . . .	209
Table 124.	Integer load byte-reverse instruction set index . . . . .	210
Table 125.	Integer load multiple instruction set index . . . . .	210
Table 126.	Integer load and reserve instruction set index . . . . .	210
Table 127.	Basic integer store instruction set index . . . . .	211
Table 128.	Integer store byte-reverse instruction set index . . . . .	211
Table 129.	Integer store multiple instruction set index . . . . .	211
Table 130.	Integer store conditional instruction set index . . . . .	211
Table 131.	Integer arithmetic instruction set index . . . . .	212
Table 132.	Integer logical instruction set index . . . . .	214
Table 133.	CR settings for compare instructions . . . . .	216
Table 134.	CR settings for integer bit test instructions . . . . .	216
Table 135.	Integer compare and bit test instruction set index . . . . .	216
Table 136.	Integer select instruction set index . . . . .	217
Table 137.	Integer trap conditions . . . . .	217
Table 138.	Integer trap instruction set index . . . . .	218
Table 139.	Integer rotate instruction set index . . . . .	218
Table 140.	Integer rotate with mask instruction set index . . . . .	219
Table 141.	Integer shift instruction set index . . . . .	219
Table 142.	Storage synchronization instruction set index . . . . .	220
Table 143.	Cache management instruction set index . . . . .	220
Table 144.	TLB management instruction set index . . . . .	220
Table 145.	Instructions listed by name . . . . .	221
Table 146.	Instructions listed by mnemonic . . . . .	228
Table 147.	List of instructions . . . . .	236
Table 148.	Interrupt types . . . . .	250
Table 149.	Interrupt registers defined by the PowerPC architecture . . . . .	253
Table 150.	Asynchronous and synchronous interrupts . . . . .	256
Table 151.	Interrupt and exception types . . . . .	259
Table 152.	Critical input interrupt register settings . . . . .	261

Table 153.	Machine check interrupt settings . . . . .	262
Table 154.	Data storage interrupt exception conditions . . . . .	263
Table 155.	Data Storage Interrupt Register Settings . . . . .	265
Table 156.	Instruction storage interrupt exception conditions . . . . .	265
Table 157.	Instruction storage interrupt register settings . . . . .	266
Table 158.	External input interrupt register settings . . . . .	266
Table 159.	Alignment interrupt register settings . . . . .	267
Table 160.	Program interrupt exception conditions . . . . .	268
Table 161.	MSR[FE0,FE1] settings . . . . .	269
Table 162.	Program interrupt register settings . . . . .	270
Table 163.	Floating-point unavailable interrupt register settings . . . . .	270
Table 164.	System call interrupt register settings . . . . .	271
Table 165.	Auxiliary processor unavailable interrupt register settings . . . . .	271
Table 166.	Decrementer interrupt register settings . . . . .	272
Table 167.	Fixed-interval timer interrupt register settings . . . . .	272
Table 168.	Watchdog timer interrupt register settings . . . . .	273
Table 169.	Data tlb error interrupt exception conditions . . . . .	273
Table 170.	Data tlb error interrupt register settings . . . . .	274
Table 171.	Instruction TLB error interrupt exception conditions . . . . .	274
Table 172.	Instruction TLB error interrupt register settings . . . . .	274
Table 173.	Debug interrupt register settings . . . . .	275
Table 174.	SPE/embedded floating-point APU unavailable interrupt register settings . . . . .	276
Table 175.	Embedded floating-point data interrupt register settings . . . . .	276
Table 176.	Embedded floating-point round interrupt register settings . . . . .	277
Table 177.	Operations to avoid . . . . .	280
Table 178.	EIS asynchronous exception priorities . . . . .	282
Table 179.	EIS synchronous exception priorities . . . . .	283
Table 180.	Load and store ordering . . . . .	295
Table 181.	Memory barrier when coherency is required (M = 1) . . . . .	295
Table 182.	Cumulative memory barrier . . . . .	296
Table 183.	Storage related MSR fields . . . . .	302
Table 184.	Exception syndrome register (ESR) definition . . . . .	303
Table 185.	Page size and EPN field comparison . . . . .	317
Table 186.	Real address generation . . . . .	318
Table 187.	Permission control for instruction, data read, and data write accesses . . . . .	319
Table 188.	Permission control and cache instructions . . . . .	320
Table 189.	TLB entry . . . . .	322
Table 190.	Fields for EA format of tlbivax . . . . .	326
Table 191.	MAS register update summary . . . . .	328
Table 192.	MAS settings for an instruction or data TLB error interrupt . . . . .	330
Table 193.	MAS settings for permissions violation ISI or DSI . . . . .	332
Table 194.	MMU assist register field updates—EIS definition . . . . .	332
Table 195.	Notation conventions . . . . .	334
Table 196.	Instruction field descriptions . . . . .	335
Table 197.	RTL notation . . . . .	337
Table 198.	Operator precedence . . . . .	340
Table 199.	Conversion models . . . . .	341
Table 200.	BI operand settings for CR fields . . . . .	364
Table 201.	BI operand settings for CR fields . . . . .	366
Table 202.	BI operand settings for CR fields . . . . .	368
Table 203.	Data samples and sizes . . . . .	371
Table 204.	Operations with special values . . . . .	586



Table 205.	Operations with special values . . . . .	591
Table 206.	Operations with special values . . . . .	593
Table 207.	Effect of SPRN[5] and MSR[PR]. . . . .	617
Table 208.	Recoding with <b>isel</b> . . . . .	668
Table 209.	Performance monitor registers—supervisor level . . . . .	669
Table 210.	Performance monitor registers—user level (read-only) . . . . .	669
Table 211.	Performance monitor apu instructions . . . . .	670
Table 212.	Mnemonic extensions for multiply accumulate instructions . . . . .	674
Table 213.	Embedded vector floating-point instruction opcodes . . . . .	677
Table 214.	Embedded scalar single-precision floating-point instruction opcodes . . . . .	678
Table 215.	Embedded scalar double-precision floating-point instruction opcodes . . . . .	678
Table 216.	EIS-defined DBSR field descriptions . . . . .	689
Table 217.	DBCR0 field descriptions . . . . .	689
Table 218.	Data storage addressing modes . . . . .	699
Table 219.	Instruction storage addressing modes . . . . .	699
Table 220.	TLB Entry 0 reset value . . . . .	702
Table 221.	MAS2 field descriptions . . . . .	703
Table 222.	MAS4 field descriptions . . . . .	703
Table 223.	System linkage instruction set index. . . . .	705
Table 224.	System register manipulation instruction set index. . . . .	705
Table 225.	Instruction Synchronization Instruction Set Index. . . . .	706
Table 226.	CR0 encodings . . . . .	707
Table 227.	Condition register setting for compare instructions. . . . .	708
Table 228.	Branch to link register instruction comparison . . . . .	708
Table 229.	Branch to count register instruction comparison. . . . .	709
Table 230.	VLE extension BO32 encodings . . . . .	709
Table 231.	VLE extension BO16 encodings . . . . .	710
Table 232.	Branch instruction set index . . . . .	710
Table 233.	Condition register instruction set index. . . . .	710
Table 234.	Basic integer load instruction set index . . . . .	711
Table 235.	Integer Load Byte-Reverse Instruction Set Index. . . . .	712
Table 236.	Integer load multiple instruction set index. . . . .	712
Table 237.	Integer load and reserve instruction set index . . . . .	712
Table 238.	Basic integer store instruction set index . . . . .	713
Table 239.	Integer store byte-reverse instruction set index . . . . .	713
Table 240.	Integer store multiple instruction set index . . . . .	713
Table 241.	Integer store conditional instruction set index. . . . .	713
Table 242.	Integer arithmetic instruction set index . . . . .	714
Table 243.	Integer logical instruction set index. . . . .	716
Table 244.	CR settings for compare instructions . . . . .	717
Table 245.	CR settings for integer bit test instructions . . . . .	718
Table 246.	Integer compare and bit test instruction set index . . . . .	718
Table 247.	Integer select instruction set index . . . . .	718
Table 248.	Integer trap conditions . . . . .	719
Table 249.	Integer trap instruction set index. . . . .	719
Table 250.	Integer rotate instruction set index . . . . .	720
Table 251.	Integer rotate with mask instruction set index. . . . .	720
Table 252.	Integer shift instruction set index . . . . .	720
Table 253.	Storage synchronization instruction set index . . . . .	721
Table 254.	Cache management instruction set index. . . . .	721
Table 255.	TLB management instruction set index . . . . .	722
Table 256.	Instructions listed by name . . . . .	723

Table 257.	Instructions listed by mnemonic . . . . .	729
Table 258.	Book E– and EIS-defined instructions listed by mnemonic . . . . .	736
Table 259.	Immediate field and displacement field encodings . . . . .	740
Table 260.	Notation conventions . . . . .	778
Table 261.	Instruction index sorted by opcode . . . . .	778
Table 262.	32-bit instruction encodings . . . . .	782
Table 263.	16-Bit VLE instructions sorted by mnemonic . . . . .	798
Table 264.	32-bit instruction encodings (by mnemonic) . . . . .	802
Table 265.	Instruction index sorted by opcode . . . . .	817
Table 266.	32-bit instruction encodings . . . . .	822
Table 267.	Instruction index sorted by mnemonic . . . . .	834
Table 268.	32-bit instructions by mnemonic (ignoring the e_ prefix) . . . . .	838
Table 269.	Instructions sorted by mnemonic (decimal and hexadecimal) . . . . .	851
Table 270.	Instructions sorted by primary opcodes (decimal and hexadecimal) . . . . .	885
Table 271.	Instructions sorted by mnemonic (binary) . . . . .	904
Table 272.	Instructions sorted by opcode (binary) . . . . .	931
Table 273.	PowerPC instruction set legend . . . . .	951
Table 274.	PowerPC instruction set legend . . . . .	956
Table 275.	Subtract immediate simplified mnemonics . . . . .	964
Table 276.	Subtract simplified mnemonics . . . . .	964
Table 277.	Word rotate and shift simplified mnemonics . . . . .	965
Table 278.	Branch instructions . . . . .	965
Table 279.	BO bit encodings . . . . .	967
Table 280.	BO operand encodings . . . . .	968
Table 281.	CR0 and CR1 fields as updated by integer instructions . . . . .	971
Table 282.	BI operand settings for CR fields for branch comparisons . . . . .	971
Table 283.	CR field identification symbols . . . . .	972
Table 284.	Branch simplified mnemonics . . . . .	973
Table 285.	Branch instructions . . . . .	973
Table 286.	Simplified mnemonics for <b>bc</b> and <b>bca</b> without LR update . . . . .	974
Table 287.	Simplified mnemonics for <b>bclr</b> and <b>bcctr</b> without LR update . . . . .	975
Table 288.	Simplified mnemonics for <b>bcl</b> and <b>bcla</b> with LR update . . . . .	975
Table 289.	Simplified mnemonics for <b>bclrl</b> and <b>bcctrl</b> with LR update . . . . .	976
Table 290.	Standard coding for branch conditions . . . . .	976
Table 291.	Branch instructions and simplified mnemonics that incorporate CR conditions . . . . .	977
Table 292.	Simplified mnemonics with comparison conditions . . . . .	977
Table 293.	Simplified mnemonics for <b>bc</b> and <b>bca</b> without comparison conditions or LR Update . . . . .	978
Table 294.	Simplified mnemonics for <b>bclr</b> and <b>bcctr</b> without comparison conditions or LR update . . . . .	979
Table 295.	Simplified mnemonics for <b>bcl</b> and <b>bcla</b> with comparison conditions, LR update . . . . .	979
Table 296.	Simplified mnemonics for <b>bclrl</b> and <b>bcctrl</b> with comparison conditions, LR update . . . . .	980
Table 297.	Word compare simplified mnemonics . . . . .	980
Table 298.	Condition register logical simplified mnemonics . . . . .	981
Table 299.	Standard codes for trap instructions . . . . .	982
Table 300.	Trap simplified mnemonics . . . . .	982
Table 301.	TO operand bit encoding . . . . .	983
Table 302.	Additional simplified mnemonics for accessing SPRGs . . . . .	984
Table 303.	Simplified mnemonics . . . . .	986
Table 304.	Shifts . . . . .	1000
Table 305.	Comparison to zero . . . . .	1004
Table 306.	Minimum and maximum . . . . .	1004
Table 307.	Simple if-then-else constructions . . . . .	1004
Table 308.	Notation conventions and general rules . . . . .	1008

---

Table 309.	Floating-point results summary—add, sub, mul, div . . . . .	1009
Table 310.	Floating-point results summary—single convert from double . . . . .	1012
Table 311.	Floating-point results summary—double convert from single . . . . .	1013
Table 312.	Floating-point results summary—convert to unsigned . . . . .	1013
Table 313.	Floating-point results summary—convert to signed . . . . .	1014
Table 314.	Floating-point results summary—convert from unsigned . . . . .	1014
Table 315.	Floating-point results summary—convert from signed . . . . .	1014
Table 316.	Floating-point results summary—*abs, *nabs, *neg . . . . .	1015
Table 317.	Glossary . . . . .	1016
Table 318.	Document revision history . . . . .	1024

## List of figures

Figure 1.	EIS programming model register set	35
Figure 2.	Effective-to-Real Address Translation Flow	41
Figure 3.	Register model	49
Figure 4.	SPE and floating point APU GPR usage	57
Figure 5.	Integer exception register (XER)	58
Figure 6.	Floating-point status and control register (FPSCR)	59
Figure 7.	Condition register (CR)	62
Figure 8.	Link register (LR)	68
Figure 9.	Count register (CTR)	69
Figure 10.	Machine state register (MSR)	70
Figure 11.	Hardware implementation dependent register 0 (HID0)	73
Figure 12.	Hardware implementation dependent register 1 (HID1)	76
Figure 13.	Processor ID register (PIR)	76
Figure 14.	Processor version register (PVR)	76
Figure 15.	System version register (SVR)	77
Figure 16.	Relationship of timer facilities to the time base.	78
Figure 17.	Timer control register (TCR)	79
Figure 18.	Timer status register (TSR)	80
Figure 19.	Time base upper/lower registers (TBU/TBL)	81
Figure 20.	Decrementer register (DEC)	82
Figure 21.	Decrementer auto-reload register (DECAR)	83
Figure 22.	Save/restore register 0 (SRR0)	83
Figure 23.	Save/restore register 1 (SRR1)	84
Figure 24.	Critical save/restore register 0 (CSRR0)	84
Figure 25.	Critical save/restore register 1 (CSRR1)	84
Figure 26.	Data exception address register (DEAR)	85
Figure 27.	Interrupt vector prefix register (IVPR)	85
Figure 28.	Interrupt vector offset registers (IVOR)	85
Figure 29.	Exception syndrome register (ESR)	87
Figure 30.	Debug save/restore register 0 (DSRR0)	89
Figure 31.	Debug save/restore register 1 (DSRR1)	89
Figure 32.	Machine check save/restore register 0 (MCSRR0)	89
Figure 33.	Machine check save/restore register 1 (MCSRR1)	90
Figure 34.	Machine check address register (MCAR/MCARU)	90
Figure 35.	Machine check syndrome register 1 (MCSR)	91
Figure 36.	Software-use sprs (SPRG0–SPRG7 and USPRG0)	92
Figure 37.	L1 cache control and status register 0 (L1CSR0)	93
Figure 38.	L1 cache control and status register 1 (L1CSR1)	95
Figure 39.	L1 cache configuration register 0 (L1CFG0)	96
Figure 40.	L1 cache configuration register 1 (L1CFG1)	97
Figure 41.	L1 flush and invalidate control register 0 (L1FINV0)	98
Figure 42.	Process ID registers (PID0–PID2)	100
Figure 43.	MMU control and status register 0 (MMUCSR0)	100
Figure 44.	MMU configuration register 1 (MMUCFG)	101
Figure 45.	TLB configuration register n (TLB0CFG)	102
Figure 46.	MAS register 0 (MAS0)	103
Figure 47.	MAS register 1 (MAS1) format	104
Figure 48.	MAS register 2 (MAS2)	105

Figure 49.	MAS register 3 (MAS3)	106
Figure 50.	MAS register 4 (MAS4)	107
Figure 51.	MAS register 5 (MAS5)	108
Figure 52.	MAS register 6 (MAS6)	108
Figure 53.	MAS register 7 (MAS7)	109
Figure 54.	Debug control register 0 (DBCR0)	110
Figure 55.	Debug control register 1 (DBCR1)	112
Figure 56.	Debug control register 2 (DBCR2)	114
Figure 57.	Debug control register 3 (DBCR3)	116
Figure 58.	Debug status register (DBSR)	117
Figure 59.	Instruction address compare registers (IAC1–IAC4)	118
Figure 60.	Data address compare registers (DAC1–DAC2)	119
Figure 61.	Data value compare registers (DVC1–DVC2)	119
Figure 62.	Signal processing, embedded floating-point status and control register (SPEFSCR)	120
Figure 63.	Alternate time base register lower (ATBL)	124
Figure 64.	Alternate time base register upper (ATBU)	124
Figure 65.	Global control register 0 (PMGC0)	126
Figure 66.	Local control A registers (PMLCa0–PMLCa3)	127
Figure 67.	Local control B registers (PMLCb0–PMLCb3)	129
Figure 68.	Performance monitor counter registers (PMC0–PMC3)	130
Figure 69.	Register indirect with immediate index addressing for integer loads/stores	159
Figure 70.	Register indirect with index addressing for integer loads/stores	160
Figure 71.	Register indirect addressing for integer loads/stores	160
Figure 72.	Branch relative addressing	166
Figure 73.	Branch conditional relative addressing	167
Figure 74.	Branch to absolute addressing	167
Figure 75.	Branch conditional to absolute addressing	168
Figure 76.	Branch conditional to link register addressing	168
Figure 77.	Branch conditional to count register addressing	169
Figure 78.	Integer and fractional operations	191
Figure 79.	Virtual Address Space in Book E	309
Figure 80.	Current address space	310
Figure 81.	Current PID Value	311
Figure 82.	Virtual address and TLB-entry comparison	313
Figure 83.	Effective-to-real address translation	313
Figure 84.	Granting of Access Permission	318
Figure 85.	TLBs accessed through MAS registers and TLB instructions	324
Figure 86.	Instruction description	351
Figure 87.	Vector absolute value ( <b>evabs</b> )	417
Figure 88.	Vector add immediate word ( <b>evaddiw</b> )	417
Figure 89.	Vector add signed, modulo, integer to accumulator word ( <b>evaddsmiaaw</b> )	418
Figure 90.	Vector add signed, saturate, integer to accumulator word ( <b>evaddssiaaw</b> )	419
Figure 91.	Vector add unsigned, modulo, integer to accumulator word ( <b>evaddumiaaw</b> )	420
Figure 92.	Vector add unsigned, saturate, integer to accumulator word ( <b>evaddusiaaw</b> )	421
Figure 93.	Vector add word ( <b>evaddw</b> )	421
Figure 94.	Vector AND ( <b>evand</b> )	422
Figure 95.	Vector AND with complement ( <b>evandc</b> )	423
Figure 96.	Vector Compare Equal ( <b>evcmpeq</b> )	424
Figure 97.	Vector compare greater than signed ( <b>evcmpgts</b> )	425
Figure 98.	Vector compare greater than unsigned ( <b>evcmpgtu</b> )	426
Figure 99.	Vector compare less than signed ( <b>evcmplt</b> )	427
Figure 100.	Vector compare less than unsigned ( <b>evcmpltu</b> )	428

Figure 101. Vector count leading signed bits word ( <b>evcntlsw</b> )	428
Figure 102. Vector count leading zeros word ( <b>evcntlzw</b> )	429
Figure 103. Vector divide word signed ( <b>evdivws</b> )	430
Figure 104. Vector divide word unsigned ( <b>evdivwu</b> )	431
Figure 105. Vector equivalent ( <b>eveqv</b> )	432
Figure 106. Vector extend sign byte ( <b>evextsb</b> )	433
Figure 107. Vector extend sign half word ( <b>evextsh</b> )	433
Figure 108. <b>evidd</b> results in big- and little-endian modes	448
Figure 109. <b>eviddx</b> results in big- and little-endian modes	448
Figure 110. <b>evldh</b> Results in Big- and Little-Endian Modes	449
Figure 111. <b>evldhx</b> results in big- and little-endian modes	450
Figure 112. <b>evldw</b> results in big- and little-endian modes	450
Figure 113. <b>evldwx</b> results in big- and little-endian modes	451
Figure 114. <b>evlhesplat</b> results in big- and little-endian modes	452
Figure 115. <b>evlhesplatx</b> results in big- and little-endian modes	453
Figure 116. <b>evlhossplat</b> results in big- and little-endian modes	453
Figure 117. <b>evlhossplatx</b> results in big- and little-endian modes	454
Figure 118. <b>evlhousplat</b> results in big- and little-endian modes	455
Figure 119. <b>evlhousplatx</b> results in big- and little-endian modes	456
Figure 120. <b>evlwe</b> results in big- and little-endian modes	456
Figure 121. <b>evlwhex</b> results in big- and little-endian modes	457
Figure 122. <b>evlwhos</b> results in big- and little-endian modes	458
Figure 123. <b>evlwhosx</b> results in big- and little-endian modes	459
Figure 124. <b>evlwhou</b> results in big- and little-endian modes	459
Figure 125. <b>evlwhoux</b> results in big- and little-endian modes	460
Figure 126. <b>evlwhsplat</b> results in big- and little-endian modes	461
Figure 127. <b>evlwhsplatx</b> results in big- and little-endian modes	462
Figure 128. <b>evlwwsplat</b> results in big- and little-endian modes	462
Figure 129. <b>evlwwsplatx</b> results in big- and little-endian modes	463
Figure 130. High order element merging ( <b>evmergehi</b> )	464
Figure 131. High order element merging ( <b>evmergehilo</b> )	464
Figure 132. Low order element merging ( <b>evmergelo</b> )	465
Figure 133. Low order element merging ( <b>evmergelohi</b> )	466
Figure 134. <b>evmhegsmfaa</b> (even form)	467
Figure 135. <b>evmhegsmfan</b> (even form)	468
Figure 136. <b>evmhegsmiaa</b> (even form)	469
Figure 137. <b>evmhegsmian</b> (even form)	470
Figure 138. <b>evmhegumiaa</b> (even form)	471
Figure 139. <b>evmhegumian</b> (even form)	472
Figure 140. Even multiply of two signed modulo fractional elements (to accumulator) ( <b>evmhesmf</b> )	473
Figure 141. Even form of vector half-word multiply ( <b>evmhesmfaaw</b> )	474
Figure 142. Even form of vector half-word multiply ( <b>evmhesmfanw</b> )	475
Figure 143. Even form for vector multiply (to accumulator) ( <b>evmhesmi</b> )	476
Figure 144. Even form of vector half-word multiply ( <b>evmhesmiaaw</b> )	477
Figure 145. Even form of vector half-word multiply ( <b>evmhesmianw</b> )	478
Figure 146. Even multiply of two signed saturate fractional elements (to accumulator) ( <b>evmhessf</b> )	479
Figure 147. Even form of vector half-word multiply ( <b>evmhessfaaw</b> )	480
Figure 148. Even form of vector half-word multiply ( <b>evmhessfanw</b> )	481
Figure 149. Even form of vector half-word multiply ( <b>evmhessiaaw</b> )	482
Figure 150. Even form of vector half-word multiply ( <b>evmhessianw</b> )	483
Figure 151. Vector multiply half words, even, unsigned, modulo, integer (to accumulator) ( <b>evmheumi</b> )	484

Figure 152. Even form of vector half-word multiply ( <b>evmheumiaaw</b> ) . . . . .	485
Figure 153. Even form of vector half-word multiply ( <b>evmheumianw</b> ) . . . . .	486
Figure 154. Even form of vector half-word multiply ( <b>evmheusiaaw</b> ) . . . . .	487
Figure 155. Even form of vector half-word multiply ( <b>evmheusianw</b> ) . . . . .	489
Figure 156. <b>evmhogsmfaa</b> (odd form) . . . . .	490
Figure 157. <b>evmhogsmfan</b> (odd form) . . . . .	491
Figure 158. <b>evmhogsmiaa</b> (odd form) . . . . .	492
Figure 159. <b>evmhogsmian</b> (odd form) . . . . .	493
Figure 160. <b>evmhogumiaa</b> (odd form) . . . . .	494
Figure 161. <b>evmhogumian</b> (odd form) . . . . .	495
Figure 162. Vector multiply half words, odd, signed, modulo, fractional (to accumulator) ( <b>evmhosmf</b> ) . . . . .	496
Figure 163. Odd form of vector half-word multiply ( <b>evmhosmfaaw</b> ) . . . . .	497
Figure 164. Odd form of vector half-word multiply ( <b>evmhosmfanw</b> ) . . . . .	498
Figure 165. Vector multiply half words, odd, signed, modulo, integer (to accumulator) ( <b>evmhosmi</b> ) . . . . .	499
Figure 166. Odd form of vector half-word multiply ( <b>evmhosmiaaw</b> ) . . . . .	500
Figure 167. Odd form of vector half-word multiply ( <b>evmhosmianw</b> ) . . . . .	501
Figure 168. Vector multiply half words, odd, signed, saturate, fractional (to accumulator) ( <b>evmhossf</b> ) . . . . .	502
Figure 169. Odd form of vector half-word multiply ( <b>evmhossfaaw</b> ) . . . . .	504
Figure 170. odd Form of Vector Half-Word Multiply ( <b>evmhossfanw</b> ) . . . . .	505
Figure 171. Odd form of vector half-word multiply ( <b>evmhossiaaw</b> ) . . . . .	507
Figure 172. Odd form of vector half-word multiply ( <b>evmhossianw</b> ) . . . . .	508
Figure 173. Vector multiply half words, odd, unsigned, modulo, integer (to accumulator) ( <b>evmhoumi</b> ) . . . . .	509
Figure 174. Odd form of vector half-word multiply ( <b>evmhoumiaaw</b> ) . . . . .	510
Figure 175. Odd form of vector half-word multiply ( <b>evmhoumianw</b> ) . . . . .	511
Figure 176. Odd form of vector half-word multiply ( <b>evmhousiaaw</b> ) . . . . .	512
Figure 177. Odd form of vector half-word multiply ( <b>evmhousianw</b> ) . . . . .	514
Figure 178. Initialize accumulator ( <b>evmra</b> ) . . . . .	514
Figure 179. Vector multiply word high signed, modulo, fractional (to accumulator) ( <b>evmwhsmf</b> ) . . . . .	515
Figure 180. Vector multiply word high signed, modulo, integer (to accumulator) ( <b>evmwhsm</b> ) . . . . .	516
Figure 181. Vector multiply word high signed, saturate, fractional (to accumulator) ( <b>evmwhssf</b> ) . . . . .	518
Figure 182. Vector multiply word high unsigned, modulo, integer (to accumulator) ( <b>evmwhumi</b> ) . . . . .	519
Figure 183. Vector multiply word low signed, modulo, integer & accumulate in words ( <b>evmwlsmiaaw</b> ) . . . . .	520
Figure 184. Vector multiply word low signed, modulo, integer and accumulate negative in words ( <b>evmwlsmlmianw</b> ) . . . . .	521
Figure 185. Vector multiply word low signed, saturate, integer & accumulate in words ( <b>evmwlsssiaaw</b> ) . . . . .	522
Figure 186. Vector multiply word low signed, saturate, integer & accumulate negative in words ( <b>evmwlsmlmsianw</b> ) . . . . .	523
Figure 187. Vector multiply word low unsigned, modulo, integer ( <b>evmwlummi</b> ) . . . . .	524
Figure 188. Vector multiply word low unsigned, modulo, integer & accumulate in words ( <b>evmwlumiaaw</b> ) . . . . .	525
Figure 189. Vector multiply word low unsigned, modulo, integer & accumulate negative in words ( <b>evmwlummlmianw</b> ) . . . . .	526
Figure 190. Vector multiply word low unsigned, saturate, integer & accumulate in words ( <b>evmwlusiaaw</b> ) . . . . .	528
Figure 191. Vector multiply word low unsigned, saturate, integer & accumulate negative in words ( <b>evmwlummlmsianw</b> ) . . . . .	529
Figure 192. Vector multiply word signed, modulo, fractional (to accumulator) ( <b>evmwsmf</b> ) . . . . .	530

Figure 193. Vector multiply word signed, modulo, fractional & accumulate ( <b>evmwsmafaa</b> ) . . . . .	531
Figure 194. Vector multiply word signed, modulo, fractional & accumulate negative ( <b>evmwsmafan</b> ) . . . . .	532
Figure 195. Vector multiply word signed, modulo, integer (to accumulator) ( <b>evmwsmafi</b> ) . . . . .	533
Figure 196. Vector multiply word signed, modulo, integer & accumulate ( <b>evmwsmaiaa</b> ) . . . . .	534
Figure 197. Vector multiply word signed, modulo, integer & accumulate negative ( <b>evmwsmaian</b> ) . . . . .	535
Figure 198. Vector multiply word signed, saturate, fractional (to accumulator) ( <b>evmwsmafa</b> ) . . . . .	536
Figure 199. Vector multiply word signed, saturate, fractional, & accumulate ( <b>evmwsmafaa</b> ) . . . . .	537
Figure 200. Vector multiply word signed, saturate, fractional & accumulate negative ( <b>evmwsmafana</b> ) . . . . .	538
Figure 201. Vector multiply word unsigned, modulo, integer (to accumulator) ( <b>evmwumi</b> ) . . . . .	539
Figure 202. Vector multiply word unsigned, modulo, integer & accumulate ( <b>evmwumiaa</b> ) . . . . .	540
Figure 203. Vector multiply word unsigned, modulo, integer & accumulate negative ( <b>evmwumian</b> ) . . . . .	541
Figure 204. Vector NAND ( <b>evnand</b> ) . . . . .	542
Figure 205. Vector negate ( <b>evneg</b> ) . . . . .	542
Figure 206. Vector NOR ( <b>evnor</b> ) . . . . .	543
Figure 207. Vector OR ( <b>evor</b> ) . . . . .	544
Figure 208. Vector OR with complement ( <b>evorc</b> ) . . . . .	545
Figure 209. Vector rotate left word ( <b>evrlw</b> ) . . . . .	546
Figure 210. Vector rotate left word immediate ( <b>evrlwi</b> ) . . . . .	546
Figure 211. Vector round word ( <b>evrndw</b> ) . . . . .	547
Figure 212. Vector select ( <b>evsel</b> ) . . . . .	548
Figure 213. Vector shift left word ( <b>evslw</b> ) . . . . .	549
Figure 214. Vector shift left word immediate ( <b>evslwi</b> ) . . . . .	549
Figure 215. Vector splat fractional immediate ( <b>evsplatfi</b> ) . . . . .	550
Figure 216. <b>evsplatfi</b> sign extend . . . . .	551
Figure 217. Vector shift right word immediate signed ( <b>evsrwis</b> ) . . . . .	551
Figure 218. Vector shift right word immediate unsigned ( <b>evsrwiu</b> ) . . . . .	552
Figure 219. Vector shift right word signed ( <b>evsrws</b> ) . . . . .	553
Figure 220. Vector shift right word unsigned ( <b>evsrwu</b> ) . . . . .	554
Figure 221. <b>evstdd</b> results in big- and little-endian modes . . . . .	554
Figure 222. <b>evstddx</b> Results in big- and little-endian modes . . . . .	555
Figure 223. <b>evstdh</b> Results in big- and little-endian modes . . . . .	556
Figure 224. <b>evstdhx</b> Results in big- and little-endian modes . . . . .	557
Figure 225. <b>evstdw</b> results in big- and little-endian modes . . . . .	558
Figure 226. <b>evstdwx</b> Results in big- and little-endian modes . . . . .	558
Figure 227. <b>evstwh</b> Results in big- and little-endian modes . . . . .	559
Figure 228. <b>evstwhex</b> Results in big- and little-endian modes . . . . .	560
Figure 229. <b>evstwho</b> Results in big- and little-endian modes . . . . .	561
Figure 230. <b>evstwhox</b> Results in big- and little-endian modes . . . . .	561
Figure 231. <b>evstwwe</b> Results in big- and little-endian modes . . . . .	562
Figure 232. <b>evstwwex</b> Results in big- and little-endian modes . . . . .	563
Figure 233. <b>evstwwo</b> Results in big- and little-endian modes . . . . .	564
Figure 234. <b>evstwwox</b> Results in big- and little-endian modes . . . . .	564
Figure 235. Vector subtract signed, modulo, integer to accumulator word ( <b>evsubsmiaaw</b> ) . . . . .	565
Figure 236. Vector subtract signed, saturate, integer to accumulator word ( <b>evsubfssiaaw</b> ) . . . . .	566
Figure 237. Vector subtract unsigned, modulo, integer to accumulator word ( <b>evsubfusiaaw</b> ) . . . . .	567
Figure 238. Vector subtract unsigned, saturate, integer to accumulator word ( <b>evsubfusiaaw</b> ) . . . . .	568
Figure 239. Vector subtract from word ( <b>evsubfw</b> ) . . . . .	569
Figure 240. Vector subtract immediate from word ( <b>evsubifw</b> ) . . . . .	570
Figure 241. Vector XOR ( <b>evxor</b> ) . . . . .	570
Figure 242. Two-element vector operations . . . . .	673
Figure 243. Floating-point data formats . . . . .	683
Figure 244. BI field (Bits 11–14 of the instruction encoding) . . . . .	970



# 1 Preface

## 1.1 About this book

The primary objective of this reference is to provide a view of the programming model defined by Book E and the Book E implementation standards (EIS).

Book E is a PowerPC™ architecture definition for embedded processors that ensures binary compatibility with the user instruction set architecture (UISA) portion of the PowerPC architecture as it was jointly developed by Apple, IBM, and Motorola (now Freescale Semiconductor, Inc.).

This book should be used with the user documentation for individual implementations; such documents provide a high-level summary of the information that appears here, as well as implementation-specific features and implementation differences that are not described here.

This document distinguishes between the three levels of the architectural and implementation definition, as follows:

- The Book E architecture —Book E defines a set of user-level instructions and registers that are drawn from the UISA portion of the AIM definition of the PowerPC architecture. Book E also include numerous other supervisor-level registers and instructions as they were defined in the AIM version of the PowerPC architecture for the virtual environment architecture (VEA) and the operating environment architecture (OEA). Because Book E defines a much different model for operating system resources such as the MMU and interrupts, it defines many new registers and instructions.
- Book E implementation standards (EIS). In many cases, the Book E architecture definition provides a very general framework, leaving many higher-level details up to the implementation.  
To ensure consistency among its Book E implementations, working standards were defined, providing an additional layer of architecture between Book E and actual devices. This layer includes more specific definitions of Book E features as well as extensions to the architecture, typically in the form of auxiliary processing units (APUs), which define additional registers, instructions, and interrupts that provide specially targeted capabilities. Note that some APUs are implementation-specific and are available only on individual devices. The APUs described here are those that are implemented on multiple processors or families of processors.  
The EIS guarantees that if an APU is implemented, it conforms to the EIS architecture described here.

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.

### 1.1.1 Audience

It is assumed that the reader has the appropriate general knowledge regarding operating systems, microprocessor system design, and the basic principles of RISC processing to use the information in this manual.

## 1.1.2 Organization

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

- [Section 1.5: Part I: Book E and Book E implementation standards](#) describes the programming model defined by the PowerPC Book E architecture and the EIS. It consists of the following chapters:
  - [Chapter 2: Overview](#), provides a general discussion of the programming, interrupt, cache, and memory management models as they are defined by Book E and the EIS.
  - [Chapter 3: Register model](#), is useful for software engineers who need to understand the programming model in general and the functionality of each register.
  - [Chapter 4: Instruction model](#), provides an overview of the addressing modes and a description of the instructions. Instructions are organized by function.
  - [Chapter 5: Interrupts and exceptions](#), provides an overview of the Book E– and EIS-defined interrupts and exception conditions that can cause them.
  - [Chapter 6: Storage architecture](#), describes the cache and MMU portions of the EIS.
  - [Chapter 7: Instruction set](#), functions as a handbook for the instruction set. Instructions are sorted by mnemonic. Each instruction description includes the instruction formats and an individualized legend that provides such information as the level or levels of the architecture in which the instruction may be found and the privilege level of the instruction.
- [Part II: EIS-defined extensions to the Book E architecture](#), describes the auxiliary procession units (APUs) defined by the EIS. It consists of the following chapters:
  - [Chapter 8: Auxiliary processing units \(APUs\)](#), describes extensions to the Book E architecture defined by the EIS. These include the following:
    - [Section 8.1: Integer select APU](#)
    - [Section 8.2: Performance monitor APU](#)
    - [Section 8.3: Signal processing engine APU \(SPE APU\)](#)
    - [Section 8.4: Embedded vector and scalar single-precision floating-point APUs \(SPFP APUs\)](#)
    - [Section 8.5: Machine check APU](#)
    - [Section 8.6: Debug APU](#)
  - [Chapter 9: Storage-related APUs](#), describes the following APUs defined by the storage architecture:
    - [Section 9.1: Cache line locking APU](#)
    - [Section 9.2: Direct cache flush APU](#)
    - [Section 9.3: Cache way partitioning APU](#)

Subsequent chapters describe the VLE extension

- [Chapter 10: VLE introduction](#)
- [Chapter 11: VLE storage addressing](#)
- [Chapter 12: VLE compatibility with the EIS](#)
- [Chapter 13: VLE instruction classes](#)
- [Chapter 14: VLE instruction set](#)
- [Chapter 15: VLE instruction index](#)

- The following appendixes are included:
  - [Appendix A: Instruction set listings](#), lists all instructions except those defined by the VLE extension instructions by both mnemonic and opcode, and includes a quick reference table with general information, such as the architecture level, privilege level, form, and whether the instruction is optional. VLE instruction opcodes are listed in [Section 14: VLE instruction set](#).
  - [Appendix B: Simplified mnemonics for PowerPC instructions](#), describes simplified mnemonics, which are provided for easier coding of assembly language programs. Simplified mnemonics are defined for the most frequently used forms of branch conditional, compare, trap, rotate and shift, and certain other instructions defined by the PowerPC™ architecture and by implementations of and extensions to the PowerPC architecture.
  - [Appendix C: Programming examples](#), gives examples of how memory synchronization instructions can be used to emulate various synchronization primitives and to provide more complex forms of synchronization. It also describes multiple precision shifts.
  - [Appendix D: Guidelines for 32-bit book E](#), provides guidelines used by 32-bit Book E implementations; a set of guidelines is also outlined for software developers. Application software written to these guidelines can be labeled 32-bit Book E applications and can be expected to execute properly on all implementations of Book E, both 32-bit and 64-bit implementations.
  - [Appendix E: Embedded floating-point results](#), provides guidelines used by 32-bit Book E implementations; a set of guidelines is also outlined for software developers. Application software written to these guidelines can be labeled 32-bit Book E applications and can be expected to execute properly on all implementations of Book E, both 32-bit and 64-bit implementations.

This book includes a glossary and an index.

## 1.2 Suggested reading

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

### 1.2.1 General information

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

- *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.

### 1.2.2 Related documentation

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

- Reference manuals—These books (formerly called user's manuals) provide details about individual implementations and are intended for use with the EREF.
- Addenda/errata to reference manuals—Because some processors have follow-on parts an addendum is provided that describes the additional features and functionality

changes. These addenda are intended for use with the corresponding reference manuals.

- Hardware specifications—Hardware specifications provide specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations.
- Technical summaries—Each device has a technical summary that provides an overview of its features. This document is roughly the equivalent to the overview (Chapter 1) of an implementation’s reference manual.
- Application notes—These short documents address specific design issues useful to programmers and engineers working with ST processors.

Additional literature is published as new processors become available.

### 1.2.3 Conventions

This document uses the following notational conventions:

**Table 1. Conventions**

Convention	Description
cleared/set	When a bit takes the value zero, it is said to be cleared; when it takes a value of one, it is said to be set.
<b>mnemonics</b>	Instruction mnemonics are shown in lowercase bold.
<i>italics</i>	Italics indicate variable command parameters, for example, <b>bcctrx</b> . Book titles in text are set in italics. Internal signals are set in italics, for example, <i>qual BG</i>
0x	Prefix to denote hexadecimal number
0b	Prefix to denote binary number
rA, rB	Instruction syntax used to identify what is typically a source GPR
rD	Instruction syntax used to identify a destination GPR
frA, frB, frC	Instruction syntax used to identify a source FPR
frD	Instruction syntax used to identify a destination FPR
REG[FIELD]	Abbreviations for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, MSR[LE] refers to the little-endian mode enable bit in the machine state register.
x	In some contexts, such as signal encodings, an unitalicized x indicates a don't care.
<i>x</i>	An italicized x indicates an alphanumeric variable.
<i>n</i>	An italicized n indicates an numeric variable.
¬	NOT logical operator
&	AND logical operator
	OR logical operator

**Table 1. Conventions (continued)**

Convention	Description
	Concatenation operator; for example, 010    111 is the same as 010111
<span style="border: 1px solid black; padding: 2px;">—</span>	Indicates a reserved field in a register. Although these bits can be written to as ones

Additional conventions used with instruction encodings are described in [Table 195](#).

## 1.3 Acronyms and abbreviations

[Table 2](#) contains acronyms and abbreviations that are used in this document.

**Table 2. Acronyms and abbreviated terms**

Term	Meaning
CR	Condition register
CTR	Count register
DTLB	Data translation lookaside buffer
EA	Effective address
ECC	Error checking and correction
FPR	Floating-point register
FPU	Floating-point unit
GPR	General-purpose register
IEEE	Institute of Electrical and Electronics Engineers
ITLB	Instruction translation lookaside buffer
L2	Secondary cache
LIFO	Last-in-first-out
LR	Link register
LRU	Least recently used
LSB	Least-significant byte
lsb	Least-significant bit
MMU	Memory management unit
MSB	Most-significant byte
msb	Most-significant bit
MSR	Machine state register
NaN	Not a number
NIA	Next instruction address
No-op	No operation

**Table 2. Acronyms and abbreviated terms (continued)**

Term	Meaning
OEA	Operating environment architecture
PTE	Page table entry
RISC	Reduced instruction set computing
RTL	Register transfer language
SIMM	Signed immediate value
SPR	Special-purpose register
TB	Time base register
TLB	Translation lookaside buffer
UIMM	Unsigned immediate value
UISA	User instruction set architecture
VA	Virtual address
VEA	Virtual environment architecture
VLE	Variable length encoding
XER	Register used primarily for indicating conditions such as carries and overflows for integer operations

## 1.4 Terminology conventions

*Table 3* lists certain terms used in this manual that differ from the architecture terminology conventions.

**Table 3. Terminology conventions**

The architecture specification	This manual
Extended mnemonics	Simplified mnemonics
Privileged mode (or privileged state)	Supervisor level
Problem mode (or problem state)	User level
Real address	Physical address
Relocation	Translation
Out-of-order memory accesses	Speculative memory accesses
Storage (locations)	Memory
Storage (the act of)	Access

*Table 4* describes instruction field notation conventions used in this manual.

Table 4. Instruction field conventions

The architecture specification	Equivalent to:
BA, BB, BT	crbA, crbB, crbD (respectively)
BF, BFA	crfD, crfS (respectively)
D	d
DS	ds
FLM	FM
FRA, FRB, FRC, FRT, FRS	frA, frB, frC, frD, frS (respectively)
$\Phi\Xi M$	CRM
RA, RB, RT, RS	rA, rB, rD, rS (respectively)
SI	SIMM
U	IMM
UI	UIMM
/, //, ///	0...0 (shaded)

## 1.5 Part I: Book E and Book E implementation standards

Part I describes the registers and instructions defined by the Book E architecture and by the Book E implementation standards (EIS). It contains the following chapters:

- [Chapter 2: Overview](#), provides a general discussion of the programming, interrupt, cache, and memory management models as they are defined by Book E and the EIS.
- [Chapter 3: Register model](#), is useful for software engineers who need to understand the programming model in general and the functionality of each register.
- [Chapter 4: Instruction model](#), provides an overview of the addressing modes and a description of the instructions. Instructions are organized by function.
- [Chapter 5: Interrupts and exceptions](#), provides an overview of the Book E– and EIS–defined interrupts and exception conditions that can cause them.
- [Chapter 6: Storage architecture](#), describes the cache and MMU portions of the EIS.

## 1.6 Part II: EIS-defined extensions to the Book E architecture

This part describes the extensions defined by the Book E Implementation Standards (EIS). It consists of the following:

- [Chapter 8: Auxiliary processing units \(APUs\)](#), describes APUs such as the **isel** instruction, performance monitor, signal processing engine (SPE), locking, and machine check APUs.
- [Chapter 9: Storage-related APUs](#), describes the following APUs defined by the storage architecture:
  - [Chapter 9.1: Cache line locking APU on page 692](#)
  - [Chapter 9.2: Direct cache flush APU on page 694](#)
  - [Chapter 9.3: Cache way partitioning APU on page 695](#)
- Subsequent chapters describe the VLE extension
  - [Chapter 10: VLE introduction](#)
  - [Chapter 11: VLE storage addressing](#)
  - [Chapter 12: VLE compatibility with the EIS](#)
  - [Chapter 13: VLE instruction classes](#)
  - [Chapter 14: VLE instruction set](#)
  - [Chapter 15: VLE instruction index](#)



## 2 Overview

This document describes the Book E version of the PowerPC™ architecture as it is further defined by the Book E implementation standards (EIS) and implemented on Book E cores.

This chapter includes overviews of the following:

- Features of the Book E version of the PowerPC architecture and implementation-details defined by the EIS
- The Book E and EIS programming model
- The Book E and EIS interrupt model
- The Book E and EIS memory management model
- Architectural compatibility and migration from the original version of the PowerPC architecture as defined by Apple, IBM, and Motorola (referred to as the AIM version of the PowerPC architecture)

### 2.1 Overview Book E and the Book E implementation standards (EIS)

Book E is a version of the PowerPC architecture intended for embedded processors. To ensure application-level compatibility with the PowerPC architecture developed by Apple, IBM and Freescale, Book E incorporates the user-level resources defined in the user instruction set architecture (UISA), Book I, of the AIM architectural definition.

Because operating systems for embedded processors have different needs than those for desktop systems, Book E defines more flexible interrupt and memory management models. Instead of the segmented memory model defined by the AIM architecture, Book E provides a page-based memory system that supports multiple variable-sized pages managed through translation lookaside buffers (TLBs). Interrupt offsets can be programmed through interrupt-specific interrupt vector offset registers (IVORs). Book E defines the interrupt vector prefix register (IVPR), which is programmed with a prefix value that is concatenated with the IVOR values to place the interrupt vector table anywhere in memory.

As a consequence, some resources defined by the AIM version of the architecture are no longer supported and new ones are provided. For example, segment and block address translation (BAT) registers are gone, and new instructions, registers, and interrupts have been defined for managing page translation and protection through TLBs.

Moreover, the Book E architecture allows greater flexibility. For example, Book E defines the TLB Write Entry (**tlbwe**) and TLB Read Entry (**tlbre**) instructions only very generally, leaving details of their execution and behavior up to the implementation. However, to ensure compatibility among Book E implementations, the Book E implementation standard (EIS) defines more specifically how these instructions work.

## 2.1.1 Auxiliary processing units (APUs)

Book E supports the use of auxiliary processing units (APUs), which allocate opcode and register space for extending the instruction set without affecting the instruction set defined by Book E. This facilitates the development of special-purpose resources that are useful to some embedded environments but impractical for others. Note that instructions from multiple APUs may be assigned the same opcode numbers of the allocated opcode space.

The EIS defines many APUs. These APUs are not required on all devices, but devices that implement them do so strictly following the EIS architectural definition. In addition, an implementation may also provide an APU that is not a part of the EIS.

APUs may consist of any combination of instructions, optional behavior of Book E–defined instructions, registers, register files, fields within Book E–defined registers, interrupts, or exception conditions within Book E–defined interrupts.

[Chapter 8: Auxiliary processing units \(APUs\)](#), provides an overview of specific APUs.

## 2.2 Instruction set

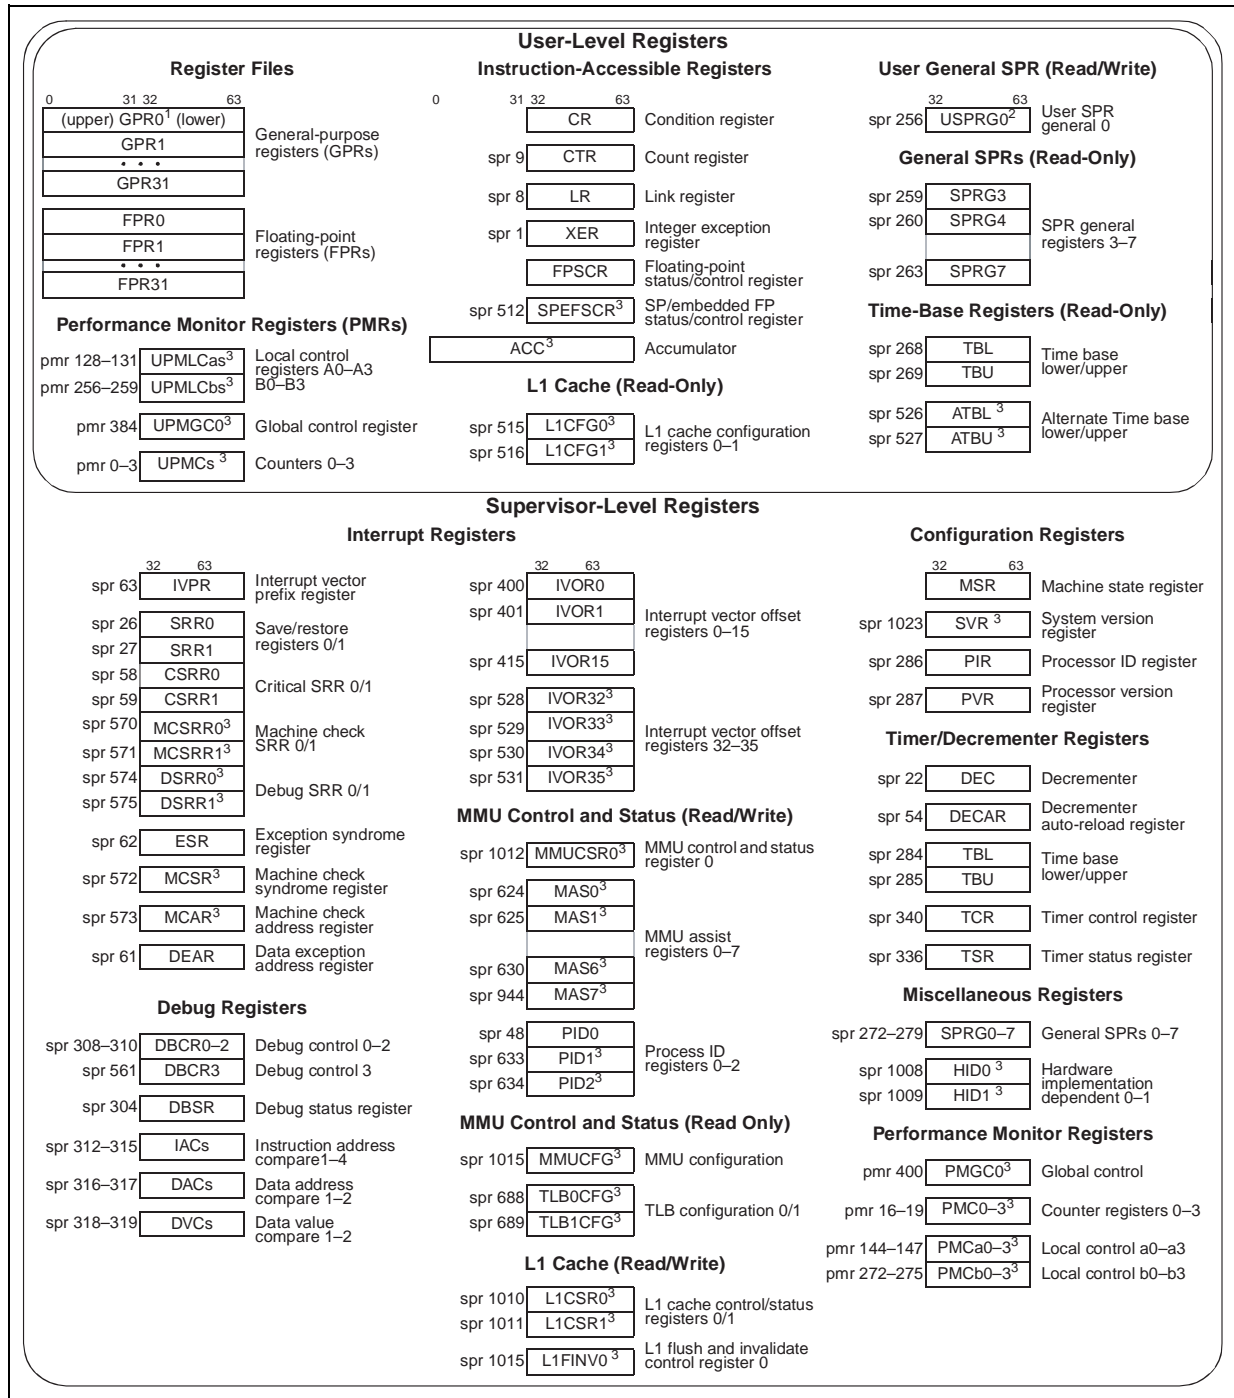
The instruction set of a ST 32-bit Book E–compliant device includes the following:

- The Book E instruction set for 32-bit implementations. This is composed primarily of the user-level instructions defined by the UISA. Some implementations do not include the Book E floating-point instructions or the Load String Word Indexed instruction (**lswx**).
- Instructions defined by EIS APUs. These include the following:
  - Integer select APU. This APU consists of the Integer Select instruction (**isel**), which incorporates an if-then-else statement that selects between two source registers by comparison to a CR bit. This instruction eliminates conditional branches, decreases branch latency, and reduces the code footprint.
  - SPE (signal processing engine) APU instructions. SPE instructions treat 64-bit GPRs as a vector of two 32-bit elements (some instructions also read or write 16-bit elements). [Section 4.6.1: SPE and embedded floating-point APUs](#), lists SPE APU vector instructions.
  - The embedded vector floating-point APU provides instructions that use the upper and lower words of the 64-bit GPRs for single-precision, vector floating-point calculations.
  - The embedded scalar single-precision APU provides instructions that use the lower 32 bits of the GPRs for single-precision, scalar floating-point calculations.
  - The embedded scalar double-precision APU instructions use the 64-bit GPRs for floating-point calculations.
  - Performance monitor APU—This APU defines two instructions, **mfpmr** and **mtpmr**, used for reading and writing the performance monitor registers (PMRs).
  - Cache block lock and unlock APU, consisting of the following instructions:
    - Data Cache Block Lock Clear (**dcblc**)
    - Data Cache Block Touch and Lock Set (**dcbtls**)
    - Data Cache Block Touch for Store and Lock Set (**dcbtstls**)
    - Instruction Cache Block Lock Clear (**icblc**)
    - Instruction Cache Block Touch and Lock Set (**icbtls**)

## 2.3 Register set

Note: Devices that implement a particular core may not implement all registers defined by that core.

Figure 1. EIS programming model register set



(1.) The 64-bit GPR registers are accessed by the SPE as separate 32-bit operands by SPE instructions. Only SPE vector instructions can access the upper word. (2.) USPRG0 is a separate physical register from SPRG0. (3.) EIS-defined registers; not part of the Book E architecture.

## 2.4 Interrupts and exception handling

Book E and the EIS support an extended exception handling model, with nested interrupt capability and extensive interrupt vector programmability. The following sections define the exception model, including an overview of exception handling as implemented in a ST Book E device, a brief description of the exception classes, and an overview of the registers involved.

### 2.4.1 Exception handling

In general, interrupt processing begins with an exception that occurs due to external conditions, errors, or program execution problems. When the exception occurs, the processor checks to verify that interrupt processing is enabled for that particular exception. If enabled, the interrupt causes the state of the processor to be saved in the appropriate registers, and prepares to begin execution of the handler located at the associated vector address for that particular exception.

Once the handler is executing, the implementation may need to check one or more bits in the exception syndrome register (ESR) or the SPEFSCR, depending on the exception type, to verify the specific cause of the exception and take appropriate action.

The interrupts are described in [Section 2.4.4: Interrupt registers](#), and in [Table 6](#).

### 2.4.2 Interrupt classes

All interrupts may be categorized as asynchronous/synchronous and critical/noncritical.

- Asynchronous interrupts are caused by events that are independent of instruction execution. The address reported in the save/restore register is that of the instruction that would have executed next had the asynchronous interrupt not occurred.
- Synchronous interrupts are caused directly by the execution or attempted execution of instructions. Synchronous inputs can be precise or imprecise:
  - Synchronous precise interrupts are those that precisely indicate the address of the instruction causing the exception that generated the interrupt or, in some cases, the address of the next instruction in program order. The interrupt type and status bits allow determination of which of the two instructions has been addressed in the appropriate save/restore register.
  - Synchronous imprecise interrupts may indicate the address of the instruction causing the exception that generated the interrupt or some instruction after the instruction causing the interrupt. If the interrupt was caused by either the context synchronizing mechanism or the execution synchronizing mechanism, the address in the appropriate save/restore register is the address of the interrupt forcing instruction. If the interrupt was not caused by either of those mechanisms, the address in the save/restore register is the last instruction to start execution and may not have completed. No instruction following the instruction in the save/restore register has executed.

### 2.4.3 Interrupt categories

Book E defines critical and noncritical interrupt categories, and the EIS defines the machine check and debug interrupt categories. Each category has a separate set of save and restore registers to which machine state and a return address are automatically written when an interrupt is taken. Each category has a return from interrupt instruction that uses the save and restore registers to reestablish the machine state of the interrupted process and

provides the address within that process at which to resume execution after the interrupt handler completes. Additional resources are provided for masking some of these interrupt categories, as described in the following:

- Debug APU interrupt (if present)—Although Book E defines debug as a critical interrupt, the EIS defines a separate debug APU. Debug save and restore registers (DSRR0/DSRR1) save state when a debug interrupt is taken; **rdci** restores state at the end of the interrupt handler. These interrupts are masked by setting the machine check enable bit, MSR[DE].
- Machine check APU interrupt (if present)—Although Book E defines machine check as a critical interrupt, the EIS defines a separate machine check APU. Machine check save and restore registers (MCSRR0/MCSRR1) save state when a machine check interrupt is taken; **rfmci** restores state at the end of the interrupt handler. These interrupts are masked by setting the machine check enable bit, MSR[ME].
- Noncritical interrupts—First-level interrupts that allow the processor to change program flow to handle conditions generated by external signals, errors, or unusual conditions arising from program execution or from programmable timer-related events. These interrupts are largely identical to those defined by the OEA portion of the Power PC architecture. They use save and restore registers (SRR0/SRR1) to save processor state and the **rfi** instruction to restore state. Asynchronous noncritical interrupts can be masked by the external interrupt enable bit, MSR[EE].
- Critical interrupts—Can be taken during a noncritical interrupt or during regular program flow. They use the critical save and restore registers (CSRR0/CSRR1) to save state when they are taken; they use the **rfci** instruction to restore state. These interrupts can be masked by the critical enable bit, MSR[CE]. Book E defines the critical input and watchdog timer interrupts as critical interrupts.

One interrupt of each category can be reported at a time; when it is taken, no program state is lost. Save/restore register pairs are serially reusable, so program state may be lost when an unordered interrupt is taken. See [Section 5.10: Interrupt ordering and masking](#).

#### 2.4.4 Interrupt registers

The registers associated with interrupt and exception handling are described in [Table 5](#).

**Table 5. Interrupt registers**

Register	Description
<b>Non critical interrupt registers</b>	
SRR0	Save/restore register 0—Stores the address of the instruction causing the exception or the address of the instruction that will execute after the <b>rfi</b> instruction.
SRR1	Save/restore register 1—Saves machine state on noncritical interrupts and restores machine state after an <b>rfi</b> instruction is executed.
<b>Critical interrupt registers</b>	
CSRR0	Critical save/restore register 0—On critical interrupts, CSRR0 stores either the address of the instruction causing the exception or the address of the instruction that will execute after the <b>rfci</b> instruction.
CSRR1	Critical save/restore register 1—CSRR1 saves machine state on critical interrupts and restores machine state after an <b>rfci</b> instruction is executed.

Table 5. Interrupt registers (continued)

Register	Description
<b>Machine check interrupt registers</b>	
MCSRR0	Machine check save/restore register 0—Stores the address of the instruction that executes after <b>rfmci</b> executes.
MCSRR1	Machine check save/restore register 1—MCSRR1 stores machine state on machine check interrupts and restores machine state (if recoverable) after an <b>rfmci</b> instruction is executed.
MCAR	Machine check address register—MCAR holds the address of the data or instruction that caused the machine check interrupt. MCAR contents are not meaningful if a signal triggered the machine check interrupt.
<b>Debug interrupt registers</b>	
DSRR0	Debug save/restore register 0—Stores the address of the instruction that executes after <b>rfdi</b> executes.
DSRR1	Debug save/restore register 1—Stores machine state on machine check interrupts and restores machine state (if recoverable) after <b>rfmci</b> executes.
<b>Syndrome registers</b>	
MCSR	Machine check syndrome register—MCSR saves machine state information on machine check interrupts and restores machine state after an <b>rfmci</b> instruction is executed.
ESR	Exception syndrome register—ESR provides a syndrome to differentiate between the different kinds of exceptions that generate the same interrupt type. Upon generation of a specific exception type, the associated bit is set and all other bits are cleared.
<b>SPE and embedded floating-point APU interrupt registers</b>	
SPEFSCR	Signal processing and embedded floating-point status and control register—Provides interrupt control and status as well as various condition bits associated with the operations performed by the SPE APU and the embedded floating-point APUs.
<b>Other interrupt registers</b>	
DEAR	Data exception address register—DEAR contains the address that was referenced by a load, store, or cache management instruction that caused an alignment, data TLB miss, or data storage interrupt.
IVPR	Interrupt vector prefix register—IVPR[32–47] contains the high-order 16 bits of the address of the exception processing routines defined in the IVOR registers.
IVORs	Interrupt vector offset registers—The IVORs contain the low-order offset of the address of the exception processing routines defined in the IVOR registers. See <a href="#">Table 6</a> .

[Table 6](#) lists IVOR registers and associated interrupts.

Table 6. Interrupt vector registers and exception conditions

Register	Interrupt
<b>Book E–defined IVORs</b>	
IVOR0	Critical input
IVOR1	Machine check interrupt offset
IVOR2	Data storage interrupt offset
IVOR3	Instruction storage interrupt offset
IVOR4	External input interrupt offset
IVOR5	Alignment interrupt offset
IVOR6	Program interrupt offset
IVOR7	Floating-point unavailable interrupt offset
IVOR8	System call interrupt offset
IVOR9	Auxiliary processor unavailable interrupt offset
IVOR10	Decrementer interrupt offset
IVOR11	Fixed-interval timer interrupt offset
IVOR12	Watchdog timer interrupt offset
IVOR13	Data TLB error interrupt offset
IVOR14	Instruction TLB error interrupt offset
IVOR15	Debug interrupt offset
<b>EIS-Defined IVORs</b>	
IVOR32	SPE APU unavailable interrupt offset
IVOR33	Embedded floating-point data exception interrupt offset
IVOR34	Embedded floating-point round exception interrupt offset
IVOR35	Performance monitor interrupt offset

Each interrupt has an associated interrupt vector address, obtained by concatenating the IVPR and IVOR values (IVPR[32–47]||IVOR $n$ [48–59]||0b0000). The resulting address is that of the instruction to be executed when that interrupt occurs. IVPR and IVOR values are indeterminate on reset, and must be initialized by the system software using **mtspr**. For more information, see [Chapter 5: Interrupts and exceptions](#).

## 2.5 Memory management

The EIS supports demand-paged virtual memory as well other memory management schemes that depend on precise control of effective-to-physical address translation and flexible memory protection as defined by Book E. The mapping mechanism consists of software-managed TLBs that support variable-sized pages with per-page properties and permissions. The following properties can be configured for each TLB:

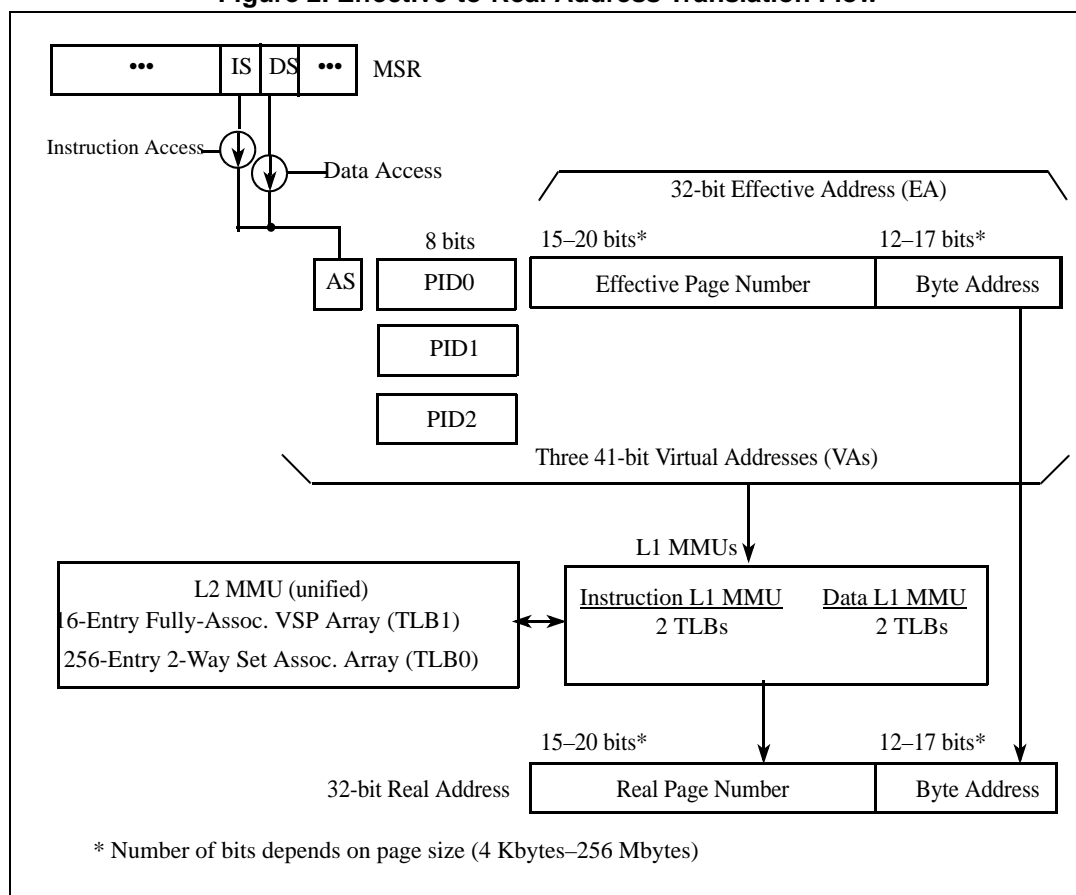
- User mode page execute access
- User mode page read access
- User mode page write access
- Supervisor mode page execute access
- Supervisor mode page read access
- Supervisor mode page write access
- Write-through required (W)
- Caching inhibited (I)
- Memory coherence required (M)
- Guarded (G)
- Endianness (E)
- User-definable (U0–U3), a 4-bit implementation-specific field

### 2.5.1 Address translation

*Figure 2* shows a typical translation flow, although each implementation may differ in the specific details. The MMU translates 32-bit effective addresses generated by loads, stores, and instruction fetches into 32-bit real addresses (used for memory bus accesses) using an interim 41-bit virtual address.



Figure 2. Effective-to-Real Address Translation Flow



As [Figure 2](#) shows, address translation starts with an effective address that is prepended with an address space (AS) value and a process ID to construct a virtual address (VA). The virtual address is then translated into a real address based on the translation information found in the on-chip TLB of the appropriate L1 MMU. The AS bit for the access is selected from the value of MSR[IS] or MSR[DS], for instruction or data accesses, respectively.

The appropriate L1 MMU (instruction or data) is checked for a matching address translation. The instruction L1 MMU and data L1 MMU operate independently and can be accessed in parallel, so that hits for instruction accesses and data accesses can occur in the same clock. If an L1 MMU misses, the request for translation is forwarded to the unified (instruction and data) L2 MMU. If found, the contents of the TLB entry are concatenated with the byte address to obtain the physical address of the requested access. On misses, the L1 TLB entries are replaced from their L2 TLB counterparts using a true LRU algorithm.

### 2.5.2 MMU assist registers (MAS1–MAS7)

Book E defines SPR numbers for the MMU assist registers, used to hold values either read from or to be written to the TLBs and information required to identify the TLB to be accessed. Book E leaves MAS register bit definitions up to the implementations. To ensure consistency among ST Book E processors, certain aspects of the implementation are defined by the Book E standard; more specific details are left to individual implementations. MAS3 implements the real page number (RPN), the user attribute bits (U0–U3), and

permission bits (UX, SX, UW, SW, UR, SR) that specify user and supervisor read, write, and execute permissions.

Some cores may not does not implement all of the MAS registers.

MAS registers are affected by the following instructions:

- MAS registers are accessed with the **mtspr** and **mfspir** instructions.
- The TLB Read Entry instruction (**tlbre**) causes the contents of a single TLB entry from the L2 MMU to be placed in defined locations in MAS0–MAS3. The TLB entry to be extracted is determined by information written to MAS0 and MAS2 before the **tlbre** instruction is executed.
- The TLB Write Entry instruction (**tlbwe**) causes the information stored in certain locations of MAS0–MAS3 to be written to the TLB specified in MAS0.
- The TLB Search Indexed instruction (**tlbsx**) updates MAS registers conditionally, based on success or failure of a lookup in the L2 MMU. The lookup is specified by the instruction encoding and specific search fields in MAS6. The values placed in the MAS registers may differ, depending on a successful or unsuccessful search.

For TLB miss and certain MMU-related DSI/ISI exceptions, MAS4 provides default values for updating MAS0–MAS2.

### 2.5.3 Process ID registers (PID0–PID2)

The Book E architecture identifies a single process ID register (PID). The EIS defines additional PIDs to hold values used to construct the virtual addresses for each access. Among these PIDs, PID0 is the Book E–defined PID. These process IDs provide an extended page sharing capability. Which of these three virtual addresses is used for translation is controlled by the TID field of a matching TLB entry, and when TID = 0x00 (identifying a page as globally shared), the PID values are ignored.

A hit to multiple TLB entries in the L1 MMU (even if they are in separate arrays) or a hit to multiple entries in the L2 MMU is considered to be a programming error.

### 2.5.4 TLB coherency

TLB entries can be invalidated as defined in the Book E architecture. The **tlbivax** instruction invalidates a matching local TLB entry.

### 2.5.5 Atomic update memory references

Book E supports atomic update memory references for both aligned word forms of data using the load and reserve and store conditional instruction pair, **lwarx** and **stwcx..** Typically, a load and reserve instruction establishes a reservation and is paired with a store conditional instruction to achieve the atomic operation. However, the programmer is responsible for preserving reservations across context switches and for protecting reservations in multiprocessor implementations.

### 2.5.6 Memory access ordering

To optimize performance, Book E supports weakly ordered references to memory. Thus, a processor manages the order and synchronization of instructions to ensure proper execution when memory is shared between multiple processes or programs. The cache and data memory control attributes, along with **msync** and **mbar**, provide the required access

control; **msync** and **mbar** are also broadcast to provide the appropriate control in the case of multiprocessor or shared memory systems.

## 2.5.7 Cache control instructions

Book E cache control instructions perform a full range of cache control functions, including cache locking by line. The EIS defines the following cache locking instructions:

- Data Cache Block Lock Clear (**dcblc**)
- Data Cache Block Touch and Lock Set (**dcbtls**)
- Data Cache Block Touch for Store and Lock Set (**dcbtstls**)
- Instruction Cache Block Lock Clear (**icblc**)
- Instruction Cache Block Touch and Lock Set (**icbtls**)

## 2.5.8 Programmable page characteristics

Cache and memory attributes are programmable on a per-page basis. In addition to the write-through, caching-inhibited, memory coherency enforce, and guarded characteristics defined by the WIMG bits, Book E defines an endianness bit, E, that selects big- or little-endian byte ordering on a per-page basis.

## 2.6 Performance monitoring

The EIS provides a performance monitoring capability that supports counting of events such as processor clocks, instruction cache misses, data cache misses, mispredicted branches, and others. The count of these events may be configured to trigger a performance monitor exception. This interrupt is assigned to vector offset register IVOR35.

The register set associated with performance monitoring consists of counter registers, a global control register, and local control registers. These registers are read/write from supervisor mode, and each register is reflected to a corresponding read-only register for user mode. The **mtpmr** and **mfpmr** instructions move data to and from these registers. An overview of the performance monitoring registers is provided in the following sections. For more information, see [Section 8.2: Performance monitor APU](#).

### 2.6.1 Global control register

The performance monitor global control register 0 (PMGC0) provides global control of the performance monitor from supervisor mode. From this register all counters may be frozen, unfrozen, or configured to freeze on an enabled condition or event. Additionally, the performance monitoring facility may be disabled or enabled from this register. The PMGC0 contents are reflected to UPMGC0, which may be read from user mode using **mfpmr**.

### 2.6.2 Performance monitor counter registers

There are four counter registers (PCM0–PCM3) provided in the performance monitor facility. These 32-bit registers hold the current count for software-selectable events and can be programmed to generate an exception on overflow. They can be accessed from supervisor mode using **mtpmr** and **mfpmr**. Their contents are reflected to UPCM0–UPCM3, which can be read from user mode with **mfpmr**.

The exception generated on overflow can be masked by clearing MSR[EE].

### 2.6.3 Local control registers

For each counter register, there are two corresponding local control registers. These two registers specify which of the 128 available events is to be counted, the action to be taken on overflow, and options for freezing a counter value under given modes or conditions.

- PMLCa0–PMLCa3 provide fields that allow freezing of the corresponding counter in user mode, supervisor mode, or under software control. The overflow condition may be enabled or disabled from these registers. Register contents are reflected to UPMCLa0–UPMLCa3, which can be read from user mode with **mfpmr**.
- PMLCb0–PMLCb3 provide count scaling for each counter register using configurable threshold and multiplier values. The threshold is a 6-bit value and the multiplier is a 3-bit encoded value, allowing 8 multiplier values in the range of 1 to 128. Any counter may be configured to increment only when an event occurs more than [threshold × multiplier] times. The contents of these registers are reflected to UPMCLb0–UPMLCb3, which can be read from user mode with **mfpmr**.

## 2.7 Legacy support of PowerPC architecture

In general, ST Book E processors support the user-level portion of the AIM architecture. The following subsections highlight the main differences. For specific details, refer to the relevant chapter.

### 2.7.1 Instruction set compatibility

The following sections generally describe compatibility between Book E and AIM PowerPC instruction sets.

#### User instruction set

The user mode instruction set defined by the AIM version of the PowerPC architecture is compatible with ST Book E processors with the following exceptions:

- Floating-point functionality provided by the embedded floating-point APUs differs from the AIM defined floating-point ISA. Also, the vector and double-precision floating-point APUs use 64-bit GPRs rather than the FPRs defined by the UISA. Most porting of floating-point operations can be handled by recompiling; however, there are new instructions specific to the APUs.
- String instructions are typically not implemented; therefore, trap emulation must be provided to ensure backward compatibility.

#### Supervisor instruction set

The supervisor mode instruction set defined by the AIM version of the PowerPC architecture is compatible with the EIS with the following exceptions:

- The MMU architecture is different, so some TLB manipulation instructions have different semantics.
- Instructions that support the BATs and segment registers are not implemented.
- Interrupt vectors are defined by the Book E  $IVOR_n$  and  $IVPR$  SPRs.
- Additional instructions are defined for returning from Book E–defined critical interrupts (**rftci**) and APU-specific interrupts.

## 2.7.2 Memory subsystem

Both Book E and the AIM version of the PowerPC architecture provide separate instruction and data memory resources. The EIS provides additional cache control features, including cache locking.

## 2.7.3 Interrupt handling

Interrupt handling is generally the same as that defined in the AIM version of the PowerPC architecture, with the following differences: (see [Section 2.4: Interrupts and exception handling](#))

- Book E defines a new critical interrupt, providing an extra level of interrupt nesting. The critical interrupt includes external critical and watchdog timer time-out inputs.
- The machine check APU implements the machine check exception differently from the Book E and from the AIM definition. It defines the Return from Machine Check Interrupt instruction, **rfmci**, and two machine check save/restore registers, MCSRR0 and MCSRR1.
- Book E processors can use IVPR and IVORs to set exception vectors individually. To provide compatibility, they can be set to the address offsets defined in the OEA.
- Unlike the AIM version of the PowerPC architecture, Book E does not define a reset vector; execution begins at a fixed virtual address, 0xFFFF\_FFFC.
- Some SPRs are different from those defined in the AIM version of the PowerPC architecture, particularly those related to the MMU functions. Much of this information has been moved to a new exception syndrome register (ESR).
- Timer services are generally compatible, although Book E defines a new decremter auto reload feature and the fixed-interval timer critical interrupt.

## 2.7.4 Memory management

ST Book E processors implement a straightforward virtual address space that complies with the Book E MMU definition, which eliminates segment registers and block address translation resources. Book E defines resources for fixed 4-Kbyte pages and multiple, variable page sizes that can be configured in a single implementation. TLB management is provided with new instructions and SPRs.

## 2.7.5 Requirements for system reset generation

Book E does not specify a system reset interrupt as was defined in the AIM version of the PowerPC architecture, but typically, system reset is initiated either by asserting a signal or by software (for example, writing a 1 to DBCR0[34], if MSR[DE] = 1

At reset, instead of invoking a reset interrupt, fetching at address 0xFFFF\_FFFC, as defined by Book E. In addition to the Book E reset definition, the EIS and the implementation define specific aspects of MMU page translation and protection mechanisms. Unlike the AIM version of the PowerPC core, as soon as instruction fetching begins, the core is in virtual mode with a hardware-initialized TLB entry.

## 2.7.6 Little-endian mode

Unlike the AIM version of the PowerPC, where the little-endian mode is controlled on a system basis, Book E supports control of byte ordering on a memory page basis. Additionally, true little-endian mode is supported by byte swapping.

## 3 Register model

This chapter describes the register model and indicates the architecture level at which each register is defined.

### 3.1 Overview

Although this chapter organizes registers according to their functionality, they can be differentiated according to how they are accessed, as follows:

- Register files. These user-level registers are accessed explicitly through source and destination operands of computational, load/store, logical, and other instructions. Book E defines two types of register files:
  - General-purpose registers (GPRs), used as source and destination operands for most operations (except Book E–defined floating-point instructions, which use FPRs). See [Section 3.3.1: General purpose registers \(GPRs\)](#).
  - Floating-point registers (FPRs), used for Book E–defined floating-point instructions. See [Section 3.4.1: Floating-point registers \(FPRs\)](#).
- Special-purpose registers (SPRs)—SPRs are accessed by using the Book E–defined Move to Special-Purpose Register (**mtspr**) and Move from Special-Purpose Register (**mfspr**) instructions. [Section 3.2.1: Special-purpose registers \(SPRs\)](#), lists SPRs.
- System-level registers that are not SPRs. These are as follows:
  - Machine state register (MSR). MSR is accessed with the Move to Machine State Register (**mtmsr**) and Move from Machine State Register (**mfmsr**) instructions. See [Section 3.6.1: Machine state register \(MSR\)](#).
  - Condition register (CR) bits are grouped into eight 4-bit fields, CR0–CR7, which are set as follows (see [Section 3.5.1: Condition register \(CR\)](#)):
    - Specified CR fields can be set by a move to the CR from a GPR (**mtcrf**).
    - A specified CR field can be set by a move to the CR from another CR field (**mcrf**), from the FPSCR (**mcrfs**), or from the XER (**mcrxr**).
    - CR0 can be set as the implicit result of an integer instruction.
    - CR1 can be set as the implicit result of a floating-point instruction.
    - A specified CR field can be set as the result of an integer or floating-point compare instruction (including SPE and SPFP compare instructions).
  - The floating-point status and control register (FPSCR). See [Section 3.4.2: Floating-point status and control register \(FPSCR\)](#).
  - The EIS-defined accumulator, which is accessed by signal processing engine (SPE) APU instructions that update the accumulator. See [Section 3.14.2: Accumulator \(ACC\)](#).
- Device control registers (DCRs). Book E defines the existence of a DCR address space and the instructions to access them, but does not define particular DCRs. The on-chip DCRs exist architecturally outside the processor core and thus are not part of Book E. The contents of DCR DCRN can be read into a GPR using **mf dcr rD,DCRN**. GPR contents can be written into DCR DCRN using **mt dcr DCRN,rS**. See [Section 3.17: Device control registers \(DCRs\)](#).
- Performance monitor registers (PMRs). (Performance monitor APU) Similar to SPRs, PMRs are accessed by using the EIS-defined Move to Performance Monitor Register (**mtpmr**) and Move from Performance Monitor Register (**mfspr**) instructions. See

[Section 3.16: Performance monitor registers \(PMRs\).](#)

## 3.2 Register model for 32-bit Book E implementations

Book E implementations include the following types of software-accessible registers:

- Registers that are accessed as part of instruction execution. These include the following:
  - The following registers are used for integer operations and are described in [Section 3.3: Registers for integer operations](#):
    - General-purpose registers (GPRs)—Book E defines a set of 32 GPRs used to hold source and destination operands for load, store, arithmetic, and computational instructions, and to read and write to other registers.
    - Integer exception register (XER)—XER bits are set based on the operation of an instruction considered as a whole, not on intermediate results. (For example, the Subtract from Carrying instruction (**subfc**), the result of which is specified as the sum of three values, sets bits in the XER based on the entire operation, not on an intermediate sum.)
  - Registers for floating-point operations. These include the following:
    - Floating-point registers (FPRs)—32 registers used to hold source and destination operands for Book E defined floating-point operations. Note that the embedded floating-point APUs do not implement FPRs; they use GPRs for floating-point operands.
    - Floating-point status and control register (FPSCR)—Used with floating-point operations. These registers are described in [Section 3.4: Registers for floating-point operations](#).
  - Condition register (CR)—Used to record conditions such as overflows and carries that occur as a result of executing arithmetic instructions (including those implemented by the SPE and SPFP APUs). The CR is described in [Section 3.5: Registers for branch operations](#).
  - Machine state register (MSR)—Used by the operating system to configure parameters such as user/supervisor mode, address space, and enabling of asynchronous interrupts. MSR is described in [Section 3.6.1: Machine state register \(MSR\)](#).
- Special-purpose registers (SPRs).
  - Book E–defined special-purpose registers (SPRs) that are accessed explicitly using **mtspr** and **mfspir** instructions. These registers are listed in [Table 7 in Section 3.2.1: Special-purpose registers \(SPRs\)](#).
  - EIS–defined SPRs that are accessed explicitly using the **mtspir** and **mfspir** instructions. These registers are listed in [Table 8 in Section 3.2.1: Special-purpose registers \(SPRs\)](#).
  - SPRs are described by function in the following sections:
    - [Section 3.5: Registers for branch operations](#)
    - [Section 3.6: Processor control registers](#)
    - [Section 3.7: Hardware implementation-dependent registers](#)
    - [Section 3.8: Timer registers](#)
    - [Section 3.9: Interrupt registers](#)
    - [Section 3.10: Software use sprs \(SPRG0–SPRG7 and USPRG0\)](#)
    - [Section 3.11: L1 cache registers](#)

- [Section 3.12: MMU registers](#)
- [Section 3.13: Debug registers](#)
- [Section 3.14: SPE and SPFP APU registers](#)
- [Section 3.15: Alternate time base registers \(ATBL and ATBU\)](#)
- EIS-defined performance monitor registers, described in [Section 3.16: Performance monitor registers \(PMRs\)](#). PMRs are like SPRs, but are accessed with EIS-defined move to and move from PMR instructions (`mtpmr` and `mfpmr`).
- EIS-defined device control registers (DCRs). Book E defines a format for implementing device-specific device-control registers. See [Section 3.17: Device control registers \(DCRs\)](#).

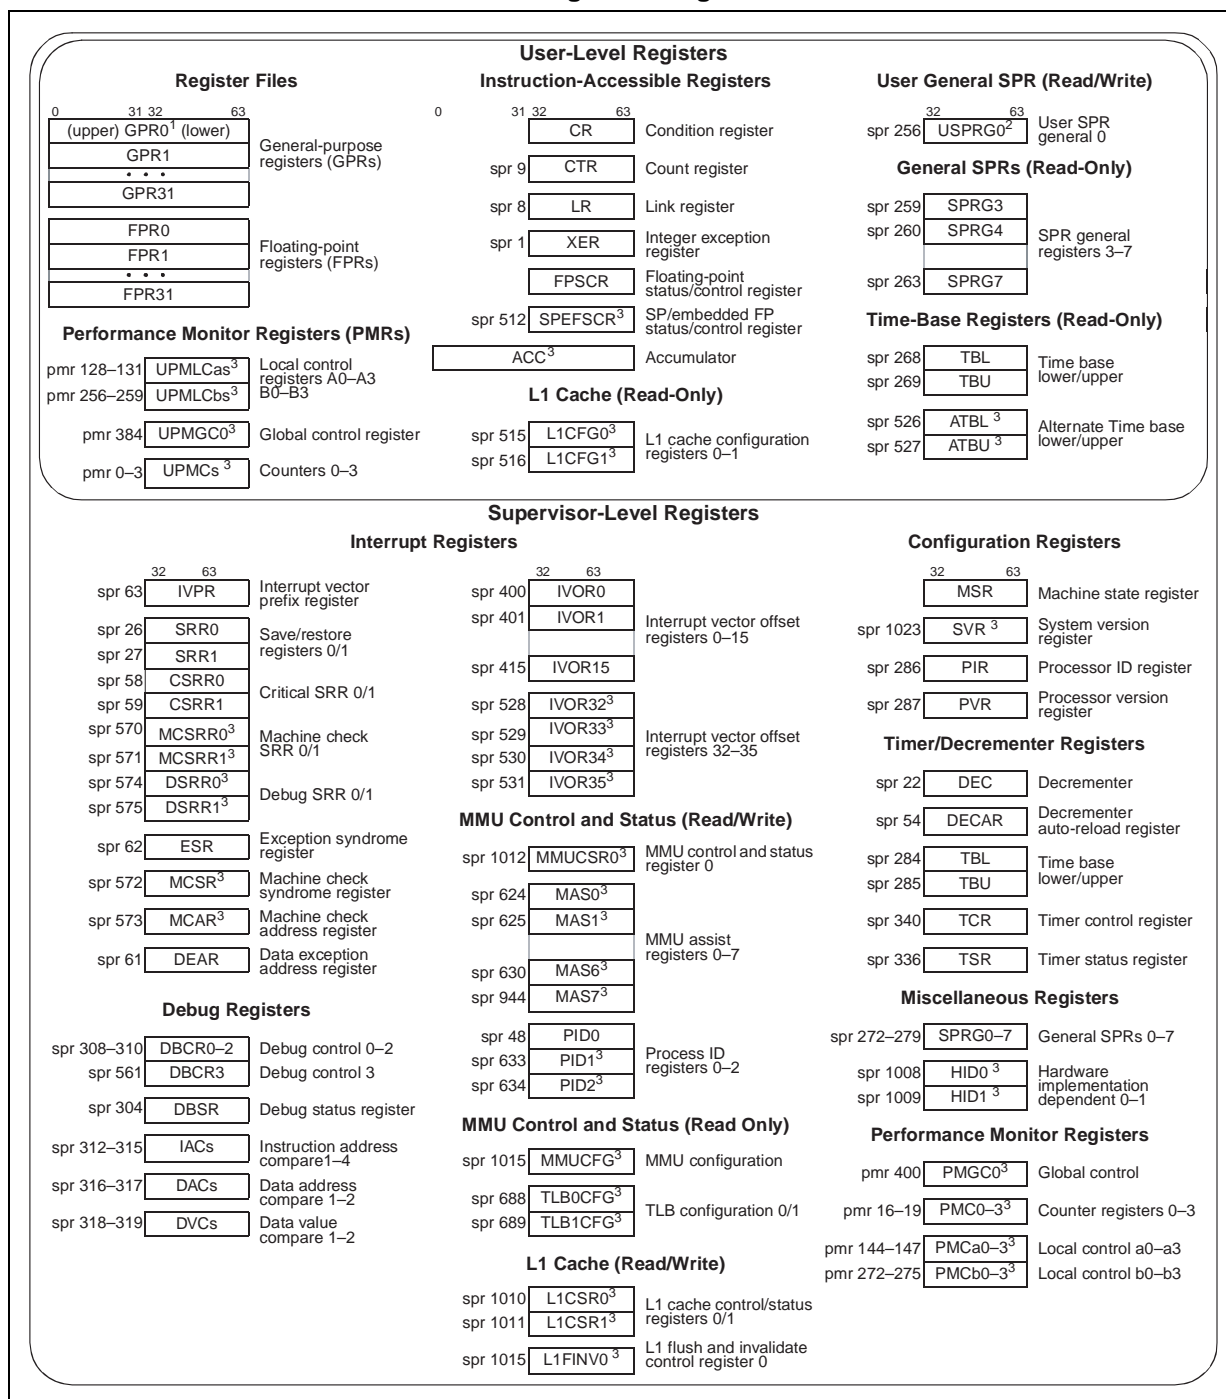
Book E defines 32- and 64-bit registers. However, except for the 64-bit FPRs, only bits 32–63 of Book E’s 64-bit registers (such as LR, CTR, the GPRs, SRR0, and CSRR0) are required to be implemented in hardware in a 32-bit Book E implementation.

Likewise, all Book E integer instructions defined to return a 64-bit result return only bits 32–63 of the result on a 32-bit Book E implementation. SPE APU vector instructions return 64-bit values; SPFP APU instructions return single-precision 32-bit values.

As with the instruction set and other aspects of the architecture, Book E defines some features very specifically, for example, resources that ensure compatibility with implementations of the PowerPC ISA. Other resources are either defined as optional or are defined in a very general way, leaving specific details up to the implementation.



Figure 3. Register model



1. The 64-bit GPR registers are accessed by the SPE as separate 32-bit operands by SPE instructions. Only SPE vector instructions can access the upper word.
2. USPRG0 is a separate physical register from SPRG0.
3. EIS-defined registers; not part of the Book E architecture.

### 3.2.1 Special-purpose registers (SPRs)

SPRs are on-chip registers that are architecturally part of the processor core. They control the use of the debug facilities, timers, interrupts, memory management unit, and other architected processor resources and are accessed with the **mtspr** and **mfspir** instructions. Unlisted encodings are reserved for future use.

*Table 7* summarizes SPRs defined in Book E. The SPR numbers are used in the instruction mnemonics. Bit 5 in an SPR number indicates whether an SPR is accessible from user or supervisor software. An **mtspr** or **mfspir** instruction that specifies an unsupported SPR number is considered an invalid instruction. Invalid instructions are treated as follows:

- If the invalid SPR falls within the range specified as user mode (SPR[5] = 0), an illegal exception is taken.
- If supervisor software attempts to access an invalid supervisor-level SPR (SPR[5] = 1), results are undefined.
- If user software attempts to access an invalid supervisor-level SPR, a privilege exception is taken.

**Table 7. Book E special purpose registers (by SPR abbreviation)**

SPR Abbreviation	Name	Defined SPR number		Access	Supervisor only
		Decimal	Binary		
CSRR0	<i>Critical save/restore register 0 (CSRR0)</i>	58	00001 11010	Read/Write	Yes
CSRR1	<i>Critical save/restore register 1 (CSRR1)</i>	59	00001 11011	Read/Write	Yes
CTR	<i>Count register (CTR)</i>	9	00000 01001	Read/Write	No
DAC1	<i>Data address compare registers (DAC1–DAC2)</i>	316	01001 11100	Read/Write	Yes
DAC2	<i>Data address compare registers (DAC1–DAC2)</i>	317	01001 11101	Read/Write	Yes
DBCR0	<i>Debug control registers (DBCR0–DBCR3) 1</i>	308	01001 10100	Read/Write	Yes
DBCR1	<i>Debug control registers (DBCR0–DBCR3) 2</i>	309	01001 10101	Read/Write	Yes
DBCR2	<i>Debug control registers (DBCR0–DBCR3) 3</i>	310	01001 10110	Read/Write	Yes
DBSR	<i>Debug status register (DBSR)</i>	304	01001 10000	Read/Clear <sup>(1)</sup>	Yes
DEAR	<i>Data exception address register (DEAR)</i>	61	00001 11101	Read/Write	Yes
DEC	<i>Decrementer register</i>	22	00000 10110	Read/Write	Yes
DECAR	<i>Decrementer auto-reload register (DECAR)</i>	54	00001 10110	Write-only	Yes
DVC1	<i>Data value compare registers (DVC1 and DVC2) 1</i>	318	01001 11110	Read/Write	Yes
DVC2	<i>Data value compare registers (DVC1 and DVC2) 2</i>	319	01001 11111	Read/Write	Yes

Table 7. Book E special purpose registers (by SPR abbreviation) (continued)

SPR Abbreviation	Name	Defined SPR number		Access	Supervisor only
		Decimal	Binary		
ESR	<i>Exception syndrome register (ESR)</i>	62	00001 11110	Read/Write	Yes
IAC1 IAC2 IAC3 IAC4	<i>Instruction address compare registers (IAC1–IAC4)</i>	312 313 314 315	01001 11000 01001 11001 01001 11010 01001 11011	Read/Write	Yes
IVOR0	<i>Interrupt vector offset registers (IVORs)</i> Critical input	400	01100 10000	Read/Write	Yes
IVOR1	<i>Interrupt vector offset registers (IVORs)</i> Machine check interrupt offset	401	01100 10001	Read/Write	Yes
IVOR10	<i>Interrupt vector offset registers (IVORs)</i> Decrementer interrupt offset	410	01100 11010	Read/Write	Yes
IVOR11	<i>Interrupt vector offset registers (IVORs)</i> Fixed-interval timer interrupt offset	411	01100 11011	Read/Write	Yes
IVOR12	<i>Interrupt vector offset registers (IVORs)</i> Watchdog timer interrupt offset	412	01100 11100	Read/Write	Yes
IVOR13	<i>Interrupt vector offset registers (IVORs)</i> Data TLB error interrupt offset	413	01100 11101	Read/Write	Yes
IVOR14	<i>Interrupt vector offset registers (IVORs)</i> Instruction TLB error interrupt offset	414	01100 11110	Read/Write	Yes
IVOR15	<i>Interrupt vector offset registers (IVORs)</i> Debug interrupt offset	415	01100 11111	Read/Write	Yes
IVOR2	<i>Interrupt vector offset registers (IVORs)</i> Data storage interrupt offset	402	01100 10010	Read/Write	Yes
IVOR3	<i>Interrupt vector offset registers (IVORs)</i> Instruction storage interrupt offset	403	01100 10011	Read/Write	Yes
IVOR4	<i>Interrupt vector offset registers (IVORs)</i> External input interrupt offset	404	01100 10100	Read/Write	Yes

Table 7. Book E special purpose registers (by SPR abbreviation) (continued)

SPR Abbreviation	Name	Defined SPR number		Access	Supervisor only
		Decimal	Binary		
IVOR5	<i>Interrupt vector offset registers (IVORs)</i> Alignment interrupt offset	405	01100 10101	Read/Write	Yes
IVOR6	<i>Interrupt vector offset registers (IVORs)</i> Program interrupt offset	406	01100 10110	Read/Write	Yes
IVOR7	<i>Interrupt vector offset registers (IVORs)</i> Floating-point unavailable interrupt offset	407	01100 10111	Read/Write	Yes
IVOR8	<i>Interrupt vector offset registers (IVORs)</i> System call interrupt offset	408	01100 11000	Read/Write	Yes
IVOR9	<i>Interrupt vector offset registers (IVORs)</i> APU unavailable interrupt offset	409	01100 11001	Read/Write	Yes
IVPR	<i>Interrupt vector offset registers (IVORs)</i> Interrupt vector	63	00001 11111	Read/Write	Yes
LR	<i>Link register (LR)</i>	8	00000 01000	Read/Write	No
PID	<i>Process ID registers (PID0–PIDn)</i>	48	00001 10000	Read/Write	Yes
PIR	<i>Processor ID register (PIR)</i>	286	01000 11110	Read-only	Yes
PVR	<i>Processor version register (PVR)</i>	287	01000 11111	Read-only	Yes
SPRG0 SPRG1 SPRG2 SPRG3 SPRG4 SPRG5 SPRG6 SPRG7	<i>Software use sprs (SPRG0–SPRG7 and USPRG0)</i>	272 273 274 275 276 277 278 279	01000 10000 01000 10001 01000 10010 01000 10011 01000 10100 01000 10101 01000 10110 01000 10111	Read/Write	Yes
SRR0	<i>Save/restore register 0 (SRR0)</i>	26	00000 11010	Read/Write	Yes
SRR1	<i>Save/restore register 1 (SRR1)</i>	27	00000 11011	Read/Write	Yes
TBL TBU	<i>Time base (TBU and TBL)</i>	284 285	01000 11100 01000 11101	Write-only	Yes
TCR	<i>Timer control register (TCR)</i>	340	01010 10100	Read/Write	Yes
TSR	<i>Timer status register (TSR)</i>	336	01010 10000	Read/Clear <sup>(2)</sup>	Yes

Table 7. Book E special purpose registers (by SPR abbreviation) (continued)

SPR Abbreviation	Name	Defined SPR number		Access	Supervisor only
		Decimal	Binary		
USPRG0	<i>Software use sprs (SPRG0–SPRG7 and USPRG0)<sup>(3)</sup></i>	256	01000 00000	Read/Write	No
USPRG3		259	01000 00011	Read-only	
USPRG4		260	01000 00100	Read-only	
USPRG5		261	01000 00101	Read-only	
USPRG6		262	01000 00110	Read-only	
USPRG7		263	01000 00111	Read-only	
UTBL		<i>Time base (TBU and TBL)</i>	268	01000 01100	
UTBU	<i>Time base (TBU and TBL)</i>	269	01000 01101	Read-only	No
XER	<i>Integer exception register (XER)</i>	1	00000 00001	Read/Write	No

1. The DBSR is read using **mfspr**. It cannot be directly written to. Instead, DBSR bits corresponding to 1 bits in the GPR can be cleared using **mtspr**.
2. The TSR is read using **mfspr**. It cannot be directly written to. Instead, TSR bits corresponding to 1 bits in the GPR can be cleared using **mtspr**.
3. User-mode read access to SPRG3 is implementation-dependent

[Table 8](#) lists EIS-defined SPRs. Compilers should recognize the mnemonic name given in this table when parsing instructions.

Table 8. EIS-defined SPRs (by SPR abbreviation)

SPR abbreviation	Name	SPR number	Access	Supervisor only	Section/page
ATBL	Alternate time base lower	526	Read-only	No	<a href="#">Section 3.15 on page 123</a>
ATBU	Alternate time base upper	527	Read-only	No	<a href="#">Section 3.15 on page 123</a>
DBCR3	Debug control register 3	561	Read/Write	Yes	<a href="#">Section 3.13.1.4 on page 116</a>
DSRR0	Debug save/restore register 0	574	R/W	Yes	<a href="#">Section 3.9.1.10 on page 89</a>
DSRR1	Debug save/restore register 1	575	R/W	Yes	<a href="#">Section 3.9.1.11 on page 89</a>
HID0	Hardware implementation dependent register 0	1008	Read/Write	Yes	<a href="#">Section 3.7.1 on page 72</a>
HID1	Hardware implementation dependent register 1	1009	Read/Write	Yes	<a href="#">Section 3.7.2 on page 75</a>
IVOR32	SPE/embedded floating-point APU unavailable interrupt offset	528	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR33	Embedded floating-point data exception interrupt offset	529	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>

Table 8. EIS–defined SPRs (by SPR abbreviation) (continued)

SPR abbreviation	Name	SPR number	Access	Supervisor only	Section/page
IVOR34	Embedded floating-point round exception interrupt offset	530	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR35	Performance monitor	531	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR36	Processor doorbell interrupt. Defined by processor signalling APU.	532	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR37	Processor doorbell critical interrupt. Defined by processor signalling APU.	533	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
L1CFG0	L1 cache configuration register 0	515	Read-only	No	<a href="#">Section 3.11.1 on page 92</a>
L1CFG1	L1 cache configuration register 1	516	Read-only	No	<a href="#">Section 3.11.2 on page 94</a>
L1CSR0	L1 cache control and status register 0	1010	Read/Write	Yes	<a href="#">Section 3.11.1 on page 92</a>
L1CSR1	L1 cache control and status register 1	1011	Read/Write	Yes	<a href="#">Section 3.11.2 on page 94</a>
L1FINV0	L1 flush and invalidate control register 0	1016	Read/Write	Yes	<a href="#">Section 3.11.5 on page 98</a>
MAS0	MMU assist register 0	624	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS1	MMU assist register 1	625	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS2	MMU assist register 2	626	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS3	MMU assist register 3	627	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS4	MMU assist register 4	628	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS5	MMU assist register 5.	629	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS6	MMU assist register 6	630	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>
MAS7	MMU assist register 7	944	Read/Write	Yes	<a href="#">Section 3.12.5 on page 103</a>

Table 8. EIS–defined SPRs (by SPR abbreviation) (continued)

SPR abbreviation	Name	SPR number	Access	Supervisor only	Section/page
MCAR	Machine check address register	573	Read-only	Yes	<a href="#">Section 3.9.1.14 on page 90</a>
MCARU	Machine check address register upper	569	Read-only	Yes	<a href="#">Section 3.9.1.14 on page 90</a>
MCSR	Machine check syndrome register	572	Read/Write	Yes	<a href="#">Section 3.9.1.15 on page 90</a>
MCSRR0	Machine-check save/restore register 0	570	Read/Write	Yes	<a href="#">Section 3.9.1.12 on page 89</a>
MCSRR1	Machine-check save/restore register 1	571	Read/Write	Yes	<a href="#">Section 3.9.1.12 on page 89</a>
MMUCFG	MMU configuration register	1015	Read-only	Yes	<a href="#">Section 3.12.3 on page 101</a>
MMUCSR0	MMU control and status register 0	1012	Read/Write	Yes	<a href="#">Section 3.12.2 on page 100</a>
PID0	Process ID register 0. Book E defines only this PID register and refers to as PID, not PID0.	48	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
PID1	Process ID register 1	633	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
PID2	Process ID register 2	634	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
SPEFSCR	Signal processing and embedded floating-point status and control register	512	Read/Write	No	<a href="#">Section 3.14.1 on page 120</a>
SVR	System version register	1023	Read-only	Yes	<a href="#">Section 3.7.5 on page 77</a>
TLB0CFG	TLB configuration register 0	688	Read-only	Yes	<a href="#">Section 3.12.4 on page 102</a>
TLB1CFG	TLB configuration register 1	689	Read-only	Yes	<a href="#">Section 3.12.4 on page 102</a>

### 3.3 Registers for integer operations

The following sections describe registers defined for integer computational instructions.

### 3.3.1 General purpose registers (GPRs)

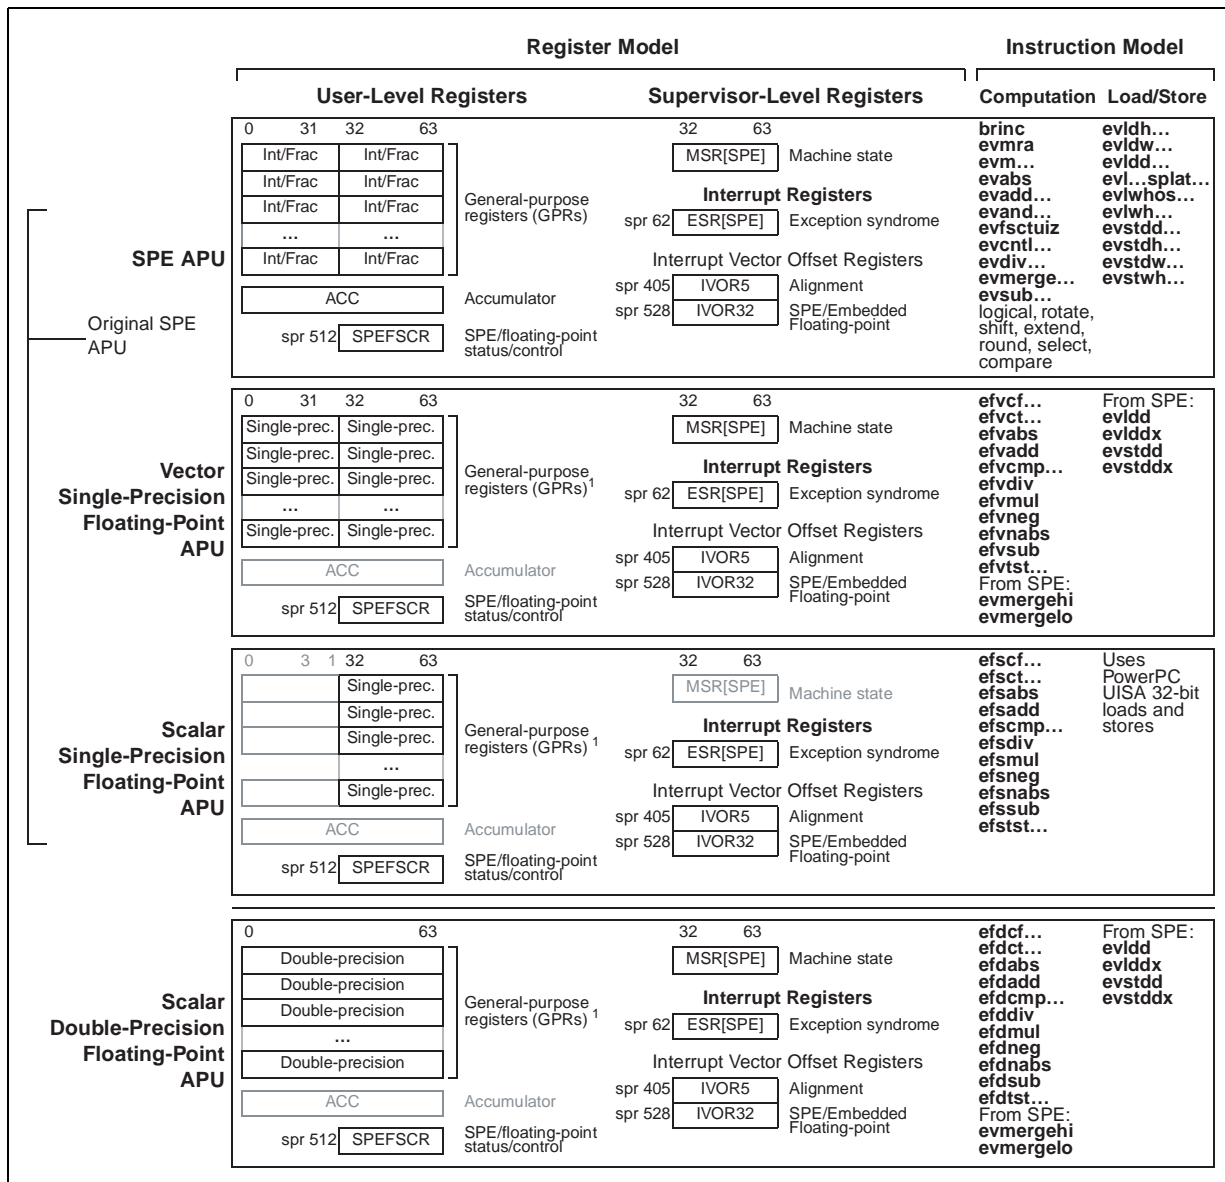
Book E implementations provide 32 GPRs (GPR0–GPR31) for integer operations. The instruction formats provide 5-bit fields for specifying the GPRs to be used in the execution of the instruction.

The Book E architecture defines 32-bit GPRs for 32-bit implementations; however, several APUs make use of GPRs that are extended to 64 bits to accommodate either vector operands or embedded double-precision floating point operands. The following APUs use the extended 64-bit GPRs:

- The signal processing engine (SPE) APU and the embedded vector single-precision floating-point APU treat the 64-bit operands as consisting of two, 32-bit elements, as shown in [Figure 4](#).
- The embedded scalar double-precision floating-point APU treats the GPRs as single 64-bit operands that accommodate IEEE double-precision values.



Figure 4. SPE and floating point APU GPR usage



1. Gray text indicates that the APU does not use this register or register field. Formatting of floating-point operands is as defined by IEEE 754, as described in the APU chapter of the EREF.

As shown in *Figure 4*, the embedded scalar single-precision floating-point APU uses 32-bit operands that, like 32-bit Book E instructions, do not affect the upper word of the 64-bit GPRs. For 32-bit implementations that implement 64-bit GPRs, all instructions except SPE APU, embedded vector single-precision APU, and embedded scalar double-precision APU instructions use and return 32-bit values in GPR bits 32–63.

### 3.3.2 Integer exception register (XER)

Bits in the integer exception register (XER) are set based on the operation of an instruction considered as a whole, not on intermediate results. (For example, the subtract from carrying

instruction (**subfc**), the result of which is specified as the sum of three values, sets bits in the XER based on the entire operation, not on an intermediate sum.)

**Figure 5. Integer exception register (XER)**




Table 9 describes XER bit definitions.

**Table 9. XER field descriptions**

Bits	Name	Description
32	SO	Summary overflow. Set when an instruction (except <b>mtspr</b> ) sets the overflow bit (OV). Once set, SO remains set until it is cleared by <b>mtspr[XER]</b> or <b>mcrxr</b> . SO is not altered by compare instructions or by other instructions (except <b>mtspr[XER]</b> and <b>mcrxr</b> ) that cannot overflow. Executing <b>mtspr[XER]</b> , supplying the values 0 for SO and 1 for OV, causes SO to be cleared and OV to be set.
33	OV	Overflow. X-form add, subtract from, and negate instructions having OE=1 set OV if the carry out of bit 32 is not equal to the carry out of bit 33, and clear OV otherwise to indicate a signed overflow. X-form multiply low word and divide word instructions having OE=1 set OV if the result cannot be represented in 32 bits ( <b>mullwo</b> , <b>divwo</b> , and <b>divwuo</b> ) and clear OV otherwise. OV is not altered by compare instructions or by other instructions (except <b>mtspr[XER]</b> and <b>mcrxr</b> ) that cannot overflow.
34	CA	Carry. Add carrying, subtract from carrying, add extended, and subtract from extended instructions set CA if there is a carry out of bit 32 and clear it otherwise. CA can be used to indicate unsigned overflow for add and subtract operations that set CA. Shift right algebraic word instructions set CA if any 1 bits are shifted out of a negative operand and clear CA otherwise. Compare instructions and instructions that cannot carry (except Shift Right Algebraic Word, <b>mtspr[XER]</b> , and <b>mcrxr</b> ) do not affect CA.
35–56	—	Reserved, should be cleared.
57–63	No. of Bytes	Supports emulation of load and store string instructions. Specifies the number of bytes to be transferred by a load string indexed or store string indexed instruction.

### 3.4 Registers for floating-point operations

This section details floating-point registers and their field descriptions.

#### 3.4.1 Floating-point registers (FPRs)

Book E defines 32 floating-point registers (FPR0–FPR31). Floating-point instruction formats provide 5-bit fields for specifying FPRs used in instruction execution.

Each FPR contains 64 bits that support the floating-point format. Instructions that interpret FPR contents as floating-point values use double-precision format for this interpretation.

The computational instructions and the move and select instructions operate on data in FPRs and, except for compare instructions, place the result into an FPR, and optionally place status information into the CR.

Load and store double instructions are provided that transfer 64 bits of data between memory and the FPRs with no conversion. Load single instructions are provided to transfer and convert floating-point values in floating-point single format from memory to the same value in floating-point double format in the FPRs. Store single instructions are provided to transfer and convert floating-point values in floating-point double format from the FPRs to the same value in floating-point single format in memory.

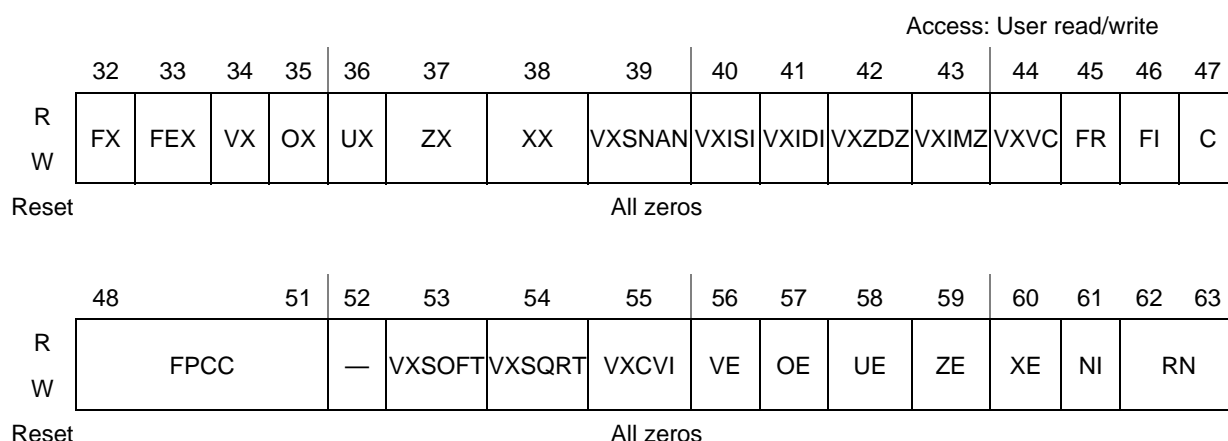
Instructions are provided that manipulate the FPSCR and the CR explicitly. Some of these instructions copy data between an FPR and the FPSCR.

The computational instructions and the select instruction accept values from the FPRs in double format. For single-precision arithmetic instructions, all input values must be representable in single format; if they are not, the result placed into the target FPR, and the setting of status bits in the FPSCR and in the CR (if Rc = 1), are undefined.

#### 3.4.2 Floating-point status and control register (FPSCR)

The FPSCR, shown below, controls how floating-point exceptions are handled and records status resulting from floating-point operations. FPSCR[32–55] are status bits; FPSCR[56–63] are control bits.

Figure 6. Floating-point status and control register (FPSCR)



The exception bits, FPSCR[35–45,53–55], are sticky; once set they remain set until they are cleared by an **mcrfs**, **mtfsfi**, **mtfsf**, or **mtfsb0**. Exception summary bits FPSCR[FX,FEX,VX] are not considered to be exception bits, and only FX is sticky.

FEX and VX are simply the ORs of other FPSCR bits, and so are not listed among the FPSCR bits affected by the various instructions. FPSCR fields are described in [Table 10](#).

**Table 10. FPSCR field descriptions**

Bits	Name	Description
32	FX	Floating-point exception summary. Every floating-point instruction, except <b>mtfsfi</b> and <b>mtfsf</b> , implicitly sets FX if that instruction causes any of the floating-point exception bits in the FPSCR to change from 0 to 1. <b>mcrfs</b> , <b>mtfsfi</b> , <b>mtfsf</b> , <b>mtfsb0</b> , and <b>mtfsb1</b> can alter FPSCR[FX] explicitly.
33	FEX	Floating-point enabled exception summary. FEX is the OR of all the floating-point exception bits masked by their respective enable bits. <b>mcrfs</b> , <b>mtfsfi</b> , <b>mtfsf</b> , <b>mtfsb0</b> , and <b>mtfsb1</b> cannot alter FPSCR[FEX] explicitly.
34	VX	Floating-point invalid operation exception summary. VX is the OR of all the invalid operation exception bits. <b>mcrfs</b> , <b>mtfsfi</b> , <b>mtfsf</b> , <b>mtfsb0</b> , and <b>mtfsb1</b> cannot alter FPSCR[VX] explicitly.
35	OX	Floating-point overflow exception
36	UX	Floating-point underflow exception
37	ZX	Floating-point zero divide exception
38	XX	Floating-point inexact exception. FPSCR[XX] is a sticky version of FPSCR[FI] (see below). Thus the following rules completely describe how FPSCR[XX] is set by a given instruction: If the instruction affects FPSCR[FI], the new FPSCR[XX] value is obtained by ORing the old value of FPSCR[XX] with the new value of FPSCR[FI]. If the instruction does not affect FPSCR[FI], the value of FPSCR[XX] is unchanged.
39	VXSNAN	Floating-point invalid operation exception (SNaN)
40	VXISI	Floating-point invalid operation exception ( $\infty - \infty$ )
41	VXIDI	Floating-point invalid operation exception ( $\infty \div \infty$ )
42	VXZDZ	Floating-point invalid operation exception ( $0 \div 0$ )
43	VXIMZ	Floating-point invalid operation exception ( $\infty \times 0$ )
44	VXVC	Floating-point invalid operation exception (invalid compare).
45	FR	Floating-point fraction rounded. The last arithmetic or rounding and conversion instruction incremented the fraction during rounding. This bit is not sticky.
46	FI	Floating-point fraction inexact. The last arithmetic or rounding and conversion instruction either produced an inexact result during rounding or caused a disabled overflow exception. This bit is not sticky. The definition of FPSCR[XX] describes the relationship between FPSCR[FI] and FPSCR[XX].
47–51	FPRF	Floating-point result flags. Set as described below in <a href="#">Table 10</a> . For arithmetic, rounding, and conversion instructions, FPRF is set based on the result placed into the target register, except that if any portion of the result is undefined, the value placed into FPRF is undefined.
47	C	Floating-point result class descriptor. Arithmetic, rounding, and conversion instructions may set this bit with the FPCC bits, to indicate the class of the result.

Table 10. FPSCR field descriptions (continued)

Bits	Name	Description
48–51	FPCC	Floating-point condition code. Floating-point Compare instructions set one of the FPCC bits and clear the other three FPCC bits. Arithmetic, rounding, and conversion instructions may set the FPCC bits with the C bit to indicate the class of the result. In this case, the three high-order FPCC bits retain their relational significance indicating that the value is less than, greater than, or equal to zero. 48 Floating-point less than or negative (FL or <) 49 Floating-point greater than or positive (FG or >) 50 Floating-point equal or zero (FE or =) 51 Floating-point unordered or NaN (FU or ?)
52	—	Reserved, should be cleared.
53	VXSOFT	Floating-point invalid operation exception (software request). Can be altered only by <b>mcrfs</b> , <b>mtfsfi</b> , <b>mtfsf</b> , <b>mtfsb0</b> , or <b>mtfsb1</b> .
54	VXSQRT	Floating-point invalid operation exception (invalid square root). Note that VXSQRT is defined even for implementations that do not support either of the two optional instructions that set it, <b>fsqrt[.]</b> and <b>frsqrt[.]</b> . Defining it for all implementations gives software a standard interface for handling square root exceptions. If an implementation does not support <b>fsqrt[.]</b> or <b>frsqrt[.]</b> , software can simulate the instruction and set VXSQRT to reflect the exception.
55	VXCVI	Floating-point invalid operation exception (invalid integer convert)
56	VE	Floating-point invalid operation exception enable
57	OE	Floating-point overflow exception enable
58	UE	Floating-point underflow exception enable
59	ZE	Floating-point zero divide exception enable
60	XE	Floating-point inexact exception enable
61	NI	Floating-point non-IEEE mode. If NI = 1, the remaining FPSCR bits may have meanings other than those given in this document and results of floating-point operations need not conform to the IEEE standard. If the IEEE-conforming result of a floating-point operation would be a denormalized number, the result of that operation is 0 (with the same sign as the denormalized number) if FPSCR[NI] = 1 and other requirements specified in the user's manual for the implementation are met. The other effects of setting NI may differ among implementations. Setting NI is intended to permit results to be approximate and to cause performance to be more predictable and less data-dependent than when NI = 0. For example, in non-IEEE mode, an implementation returns 0 instead of a denormalized number and may return a large number instead of an infinity. In non-IEEE mode an implementation should provide a means for ensuring that all results are produced without software assistance (that is, without causing an enabled exception type program interrupt or a floating-point unimplemented instruction exception type program interrupt and without invoking an emulation assist). The means may be controlled by one or more other FPSCR bits (recall that the other FPSCR bits have implementation-dependent meanings if NI = 1).
62–63	RN	Floating-point rounding control (RN). 00 Round to nearest 01 Round toward zero 10 Round toward +infinity 11 Round toward –infinity

Table 11 describes floating-point result flags.

**Table 11. Floating-point result flags**

Result flags					Result value class
C	<	>	=	?	
1	0	0	0	1	Quiet NaN
0	1	0	0	1	–Infinity
0	1	0	0	0	–Normalized number
1	1	0	0	0	–Denormalized number
1	0	0	1	0	–Zero
0	0	0	1	0	+Zero
1	0	1	0	0	+Denormalized number
0	0	1	0	0	+Normalized number
0	0	1	0	1	+Infinity


### 3.5 Registers for branch operations

This section describes registers used by Book E branch and CR operations.

#### 3.5.1 Condition register (CR)

The 32-bit CR reflects the result of certain operations and provides a mechanism for testing and branching.

**Figure 7. Condition register (CR)**



CR bits are grouped into eight 4-bit fields, CR0–CR7, which are set as follows:

- Specified CR fields can be set by a move to the CR from a GPR (**mtrcf**).
- A specified CR field can be set by a move to the CR from another CR field (**mcrf**), from the FPSCR (**mcrfs**), or from the XER (**mcrxr**).
- CR0 can be set as the implicit result of an integer instruction.
- CR1 can be set as the implicit result of a floating-point instruction.
- A specified CR field can be set as the result of either an integer or a floating-point compare instruction (including SPE and SPFP compare instructions).

Instructions are provided to perform logical operations on individual CR bits and to test individual CR bits (see [Section 4.6.6.4: Condition register instructions](#)).

Note that instructions that access CR bits (for example, Branch Conditional (**bc**), CR logicals, and Move to Condition Register Field (**mtcrf**)) determine the bit position by adding 32 to the operand value. For example, in conditional branch instructions, the BI operand accesses bit BI + 32, as shown in [Table 12](#).

**Table 12. BI operand settings for CR fields**

CRn Bits	CR Bits	BI	Description
CR0[0]	32	00000	Negative (LT)—Set when the result is negative. For SPE compare and test instructions: Set if the high-order element of rA is equal to the high-order element of rB; cleared otherwise.
CR0[1]	33	00001	Positive (GT)—Set when the result is positive (and not zero). For SPE compare and test instructions: Set if the low-order element of rA is equal to the low-order element of rB; cleared otherwise.
CR0[2]	34	00010	Zero (EQ)—Set when the result is zero. For SPE compare and test instructions: Set to the OR of the result of the compare of the high and low elements.
CR0[3]	35	00011	Summary overflow (SO). Copy of XER[SO] at the instruction's completion. For SPE compare and test instructions: Set to the AND of the result of the compare of the high and low elements.
CR1[0]	36	00100	Copy of FPSCR[FX] at the instruction's completion. Negative (LT) For SPE and SPFP compare and test instructions: Set if the high-order element of rA is equal to the high-order element of rB; cleared otherwise.
CR1[1]	37	00101	Copy of FPSCR[FEX] at the instruction's completion. Positive (GT) For SPE and SPFP compare and test instructions: Set if the low-order element of rA is equal to the low-order element of rB; cleared otherwise.
CR1[2]	38	00110	Copy of FPSCR[VX] at the instruction's completion. Zero (EQ) For SPE and SPFP compare and test instructions: Set to the OR of the result of the compare of the high and low elements.
CR1[3]	39	00111	Copy of FPSCR[OX] at the instruction's completion. Summary overflow (SO) For SPE and SPFP compare and test instructions: Set to the AND of the result of the compare of the high and low elements.
CRn[0]	40 44 48 52 56 60	01000 01100 10000 10100 11000 11100	Less than or floating-point less than (LT, FL). For integer compare instructions: rA < SIMM or rB (signed comparison) or rA < UIMM or rB (unsigned comparison). For floating-point compare instructions: frA < frB. For SPE and SPFP compare and test instructions: Set if the high-order element of rA is equal to the high-order element of rB; cleared otherwise.

Table 12. BI operand settings for CR fields (continued)

CRn Bits	CR Bits	BI	Description
CRn[1]	41	01001	Greater than or floating-point greater than (GT, FG).
	45	01101	For integer compare instructions:
	49	10001	rA > SIMM or rB (signed comparison) or rA > UIMM or rB (unsigned comparison).
	53	10101	For floating-point compare instructions: frA > frB.
	57	11001	For SPE and SPFP compare and test instructions:
	61	11101	Set if the low-order element of rA is equal to the low-order element of rB; cleared otherwise.
CRn[2]	42	01010	Equal or floating-point equal (EQ, FE).
	46	01110	For integer compare instructions: rA = SIMM, UIMM, or rB.
	50	10010	For floating-point compare instructions: frA = frB.
	54	10110	For SPE and SPFP compare and test instructions:
	58	11010	Set to the OR of the result of the compare of the high and low elements.
	62	11110	
CRn[3]	43	01011	Summary overflow or floating-point unordered (SO, FU).
	47	01111	For integer compare instructions, this is a copy of XER[SO] at the completion of the instruction.
	51	10011	
	55	10111	For floating-point compare instructions, one or both of frA and frB is a NaN.
	59	11011	For SPE and SPFP vector compare and test instructions:
	63	11111	Set to the AND of the result of the compare of the high and low elements.

### 3.5.1.1 CR setting for integer instructions

For all integer word instructions in which the Rc bit is defined and set, and for **addic.**, **andi.**, and **andis.**, CR0[32–34] are set by signed comparison of bits 32–63 of the result to zero; CR[35] is copied from the final state of XER[SO]. The Rc bit is not defined for double-word integer operations.

```

if      (target_register)32-63 < 0 then c ← 0b100
else if (target_register)32-63 > 0 then c ← 0b010
else                                     c ← 0b001
CR0 ← c || XERSO

```

The value of any undefined portion of the result is undefined, and the value placed into the first three bits of CR0 is undefined. CR0 bits are interpreted as described in [Table 13](#).

Table 13. CR0 bit descriptions

CR bit	Name	Description
32	Negative (LT)	Bit 32 of the result is equal to one.
33	Positive (GT)	Bit 32 of the result is equal to zero, and at least one of bits 33–63 of the result is non-zero.
34	Zero (EQ)	Bits 32–63 of the result are equal to zero.
35	Summary overflow (SO)	This is a copy of the final state of XER[SO] at the completion of the instruction.



Note that CR0 may not reflect the true (infinitely precise) result if overflow occurs.

**3.5.1.2 CR setting for store conditional instructions**

CR0 is also set by the integer store conditional instruction, **stwcx.**. See instruction descriptions in [Chapter 4: Instruction model](#), for detailed descriptions of how CR0 is set.

**3.5.1.3 CR setting for floating-point instructions**

For all floating-point instructions in which the Rc bit is defined and set, CR1 (CR[36–39]) is copied from FPSCR[32–35]. These bits are interpreted as shown in [Table 14](#).

**Table 14. CR setting for floating-point instructions**

Bit	Name	Description
36	FX	Floating-point exception summary. Copy of final state of FPSCR[FX] at instruction completion.
37	FEX	Floating-point enabled exception summary. Copy of final state of FPSCR[FEX] at instruction completion.
38	VX	Floating-point invalid operation exception summary. Copy of final state of FPSCR[VX] at completion.
39	OX	Floating-point overflow exception. Copy of final state of FPSCR[OX] at instruction completion.

**3.5.1.4 CR setting for compare instructions**

For compare instructions, a CR field specified by the BI field in the instruction is set to reflect the result of the comparison, as shown in [Table 15](#).

**Table 15. CR setting for compare instructions**

CRn bit	Bit expression	CR Bits		BI		Description
		AIM (BI Operand)	Book E	0–2	3–4	
CRn[0]	4 * cr0 + lt (or lt)	0	32	000	00	Less than or floating-point less than (LT, FL). For integer compare instructions: rA < SIMM or rB (signed comparison) or rA < UIMM or rB (unsigned comparison). For floating-point compare instructions: frA < frB.
	4 * cr1 + lt	4	36	001		
	4 * cr2 + lt	8	40	010		
	4 * cr3 + lt	12	44	011		
	4 * cr4 + lt	16	48	100		
	4 * cr5 + lt	20	52	101		
	4 * cr6 + lt	24	56	110		
	4 * cr7 + lt	28	60	111		

Table 15. CR setting for compare instructions (continued)

CRn bit	Bit expression	CR Bits		BI		Description
		AIM (BI Operand)	Book E	0–2	3–4	
CRn[1]	4 * cr0 + gt (or gt)	1	33	000	01	Greater than or floating-point greater than (GT, FG). For integer compare instructions: rA > SIMM or rB (signed comparison) or rA > UIMM or rB (unsigned comparison). For floating-point compare instructions: frA > frB.
	4 * cr1 + gt	5	37	001		
	4 * cr2 + gt	9	41	010		
	4 * cr3 + gt	13	45	011		
	4 * cr4 + gt	17	49	100		
	4 * cr5 + gt	21	53	101		
	4 * cr6 + gt	25	57	110		
4 * cr7 + gt	29	61	111			
CRn[2]	4 * cr0 + eq (or eq)	2	34	000	10	Equal or floating-point equal (EQ, FE). For integer compare instructions: rA = SIMM, UIMM, or rB. For floating-point compare instructions: frA = frB.
	4 * cr1 + eq	6	38	001		
	4 * cr2 + eq	10	42	010		
	4 * cr3 + eq	14	46	011		
	4 * cr4 + eq	18	50	100		
	4 * cr5 + eq	22	54	101		
	4 * cr6 + eq	26	58	110		
4 * cr7 + eq	30	62	111			
CRn[3]	4 * cr0 + so/un (or so/un)	3	35	000	11	Summary overflow or floating-point unordered (SO, FU). For integer compare instructions, this is a copy of XER[SO] at instruction completion. For floating-point compare instructions, one or both of frA and frB is a NaN.
	4 * cr1 + so/un	7	39	001		
	4 * cr2 + so/un	11	43	010		
	4 * cr3 + so/un	15	47	011		
	4 * cr4 + so/un	19	51	100		
	4 * cr5 + so/un	23	55	101		
	4 * cr6 + so/un	27	59	110		
4 * cr7 + so/un	31	63	111			

### 3.5.1.5 CR bit settings in VLE mode

The VLE extension implements the entire CR, but some comparison operations and all branch instructions are limited to using CR0–CR3. However, all Book E CR field and logical operations are provided.

CR bits are grouped into eight 4-bit fields, CR0–CR7, which are set in one of the following ways.

- Specified CR fields can be set by a move to the CR from a GPR (**mtcrf**).
- A specified CR field can be set by a move to the CR from another CR field (**e\_mcrf**).
- CR field 0 can be set as the implicit result of an integer instruction.
- A specified CR field can be set as the result of an integer compare instruction.
- CR field 0 can be set as the result of an integer bit test instruction.

Instructions are provided to perform logical operations on individual CR bits and to test individual CR bits.

CR settings for integer instructions

For all integer word instructions in which the Rc bit is defined and set, and for **addic.**, the first three bits of CR field 0 (CR[32–34]) are set by signed comparison of bits 32–63 of the result to zero, and the fourth bit of CR field 0 (CR[35]) is copied from the final state of XER[SO].

```

if      (target_register)32:63 < 0 then c ← 0b100
else if (target_register)32:63 > 0 then c ← 0b010
else                                     c ← 0b001
CR0 ← c || XERSO
    
```

If any portion of the result is undefined, the value placed into the first three bits of CR field 0 is undefined. The bits of CR field 0 are interpreted as shown in [Table 16](#).

**Table 16. CR0 encodings**

CR bit	Description
32	Negative (LT). Bit 32 of the result is equal to 1.
33	Positive (GT). Bit 32 of the result is equal to 0 and at least one of bits 33–63 of the result is non-zero.
34	Zero (EQ). Bits 32–63 of the result are equal to 0.
35	Summary overflow (SO). This is a copy of the final state XER[SO] at the completion of the instruction.

CR setting for compare instructions supported by the VLE extension

For compare instructions, a CR field specified by the **crD** operand in the instruction for the **e\_cmp**, **e\_cmpl**, **e\_cmphi**, and **e\_cmpli** instructions, or CR0 for the **e\_cmp16i**, **e\_cmpl16i**, **e\_cmphi16i**, **e\_cmpli16i**, **se\_cmp**, **se\_cmpl**, **se\_cmphi**, **se\_cmpli**, and **se\_cmpli** instructions is set to reflect the result of the comparison. The CR field bits are interpreted as shown in [Table 17](#). A complete description of how the bits are set is given in [Chapter 7: Instruction set](#), and in [Section 4.6.6.5: Integer instructions](#).

**Table 17. Condition register setting for compare instructions**

CR bit	Description
4×CRD + 32	Less than (LT). For signed-integer compare, GPR(rA or rX) < SCI8 or SI or GPR(rB or rY). For unsigned-integer compare, GPR(rA or rX) < <sub>u</sub> SCI8 or UI or UI5 or GPR(rB or rY).
4×CRD + 33	Greater than (GT). For signed-integer compare, GPR(rA or rX) > SCI8 or SI or UI5 or GPR(rB or rY). For unsigned-integer compare, GPR(rA or rX) > <sub>u</sub> SCI8 or UI or UI5 or GPR(rB or rY).
4×CRD + 34	Equal (EQ). For integer compare, GPR(rA or rX) = SCI8 or UI5 or SI or UI or GPR(rB or rY).
4×CRD + 35	Summary overflow (SO). For integer compare, this is a copy of the final state of XER[SO] at the completion of the instruction.


### 3.5.1.6 CR setting for the VLE bit test instruction

The Bit Test Immediate instruction, **se\_btsti**, also sets CR field 0. See the instruction description and also [Section 4.6.6.5: Integer instructions](#).

### 3.5.2 Link register (LR)

The link register can be used to provide the branch target address for a Branch Conditional to LR (**bclrx**) instruction, and it holds the return address after branch and link instructions.

**Figure 8. Link register (LR)**



The LR contents are read into a GPR using **mf spr**. The contents of a GPR can be written to the LR using **mt spr**. LR[62–63] are ignored by **bclr** instructions.

#### 3.5.2.1 Link register usage in VLE mode

VLE instructions use the LR as defined in Book E, although the VLE extension defines a subset of all variants of Book E conditional branches involving the LR, as shown in [Table 18](#). Note that because VLE instructions can reside on half-word boundaries, in VLE mode, LR[30] is examined when the LR holds an instruction address.

**Table 18. Branch to link register instruction comparison**


Book E		VLE Subset	
Instruction	Syntax	Instruction	Syntax
Branch Conditional to Link Register	<b>bclr</b> BO,BI	Branch (Absolute) to Link Register	<b>se_blr</b>
Branch Conditional to Link Register & Link	<b>bclrl</b> BO,BI	Branch (Absolute) to Link Register & Link	<b>se_blrl</b>
Branch Conditional & Link	<b>e_bcl</b> BO,BI,BD	Branch Conditional & Link	<b>e_bcl</b> BO32,BI32,BD15
		Branch (Absolute) & Link	<b>e_bl</b> BD24 <b>se_bl</b> BD8

### 3.5.3 Count register (CTR)

CTR can be used to hold a loop count that can be decremented and tested during execution of branch instructions that contain an appropriately encoded BO field. If the CTR value is 0 before being decremented, it is –1 afterward. The entire CTR can be used to hold the branch target address for a Branch Conditional to CTR (**bcctrx**) instruction.

Note that because VLE instructions can reside on half-word boundaries, in VLE mode, CTR[30] is examined when the CTR holds an instruction address.

Figure 9. Count register (CTR)



### 3.5.3.1 Count register usage in VLE mode

VLE instructions use the CTR as defined by in Book E, although the VLE extension defines a subset of the variants of Book E conditional branches involving the CTR, as shown in [Table 19](#).

Table 19. Branch to count register instruction comparison

Book E		VLE	
Instruction	Syntax	Instruction	Syntax
Branch conditional to count register Branch conditional to count register & link	<b>bcctr</b> BO,BI <b>bcctrl</b> BO,BI	Branch (absolute) to count register Branch (absolute) to count register & link	<b>se_bctr</b> <b>se_bctrl</b>

## 3.6 Processor control registers

This section addresses machine state, processor ID, and processor version registers.

### 3.6.1 Machine state register (MSR)

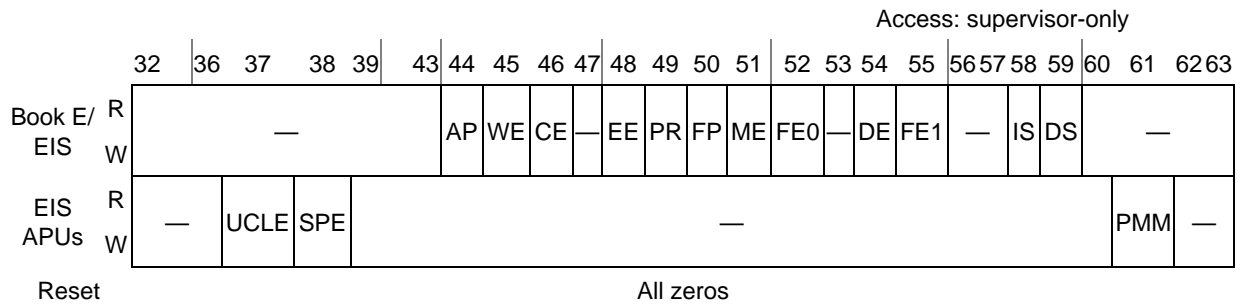
The MSR defines the state of the processor (that is, enabling and disabling of interrupts and debugging exceptions, enabling and disabling of address translation for instruction and data memory accesses, enabling and disabling some APUs, and specifying whether the processor is in supervisor or user mode).

MSR contents are automatically saved, altered, and restored by the interrupt-handling mechanism. If a non-critical interrupt is taken, MSR contents are automatically copied into SRR1. If a critical interrupt is taken, MSR contents are automatically copied into CSRR1. When an **rfi** or **rfci** is executed, MSR contents are restored from SRR1 or CSRR1.

The EIS-defined machine check APU defines additional save/restore resources. When a machine check interrupt is taken, MCSRR0 and MCSRR1 hold the return address and MSR information. The return from machine check interrupt instruction, **rfmci**, restores MCSRR1 contents to the MSR.

MSR contents are read into a GPR using **mfmsr**. The contents of a GPR can be written to MSR using **mtmsr**. The write MSR external enable instructions (**wrtee** and **wrteei**) can be used to set or clear MSR[EE] without affecting other MSR bits.

**Figure 10. Machine state register (MSR)**



**Table 20. MSR field descriptions**

Bits	Name	Description
32–36	—	Reserved, should be cleared. <sup>(1)</sup>
37	UCLE	(Cache-locking APU) User-mode cache lock enable. Used to restrict user-mode cache-line locking by the operating system. 0 Any cache lock instruction executed in user-mode takes a cache-locking DSI exception and sets either ESR[DLK] or ESR[ILK]. This allows the operating system to manage and track the locking/unlocking of cache lines by user-mode tasks. 1 Cache-locking instructions can be executed in user-mode and they do not take a DSI for cache-locking. (They may still take a DSI for access violations though.)
38	SPE	(SPE, SPFP, DPFP APUs) SPE enable. Enables use of 64-bit extended GPRs used by SPE, single-precision vector, and double-precision floating-point APUs/ 0 If software attempts to execute an SPE APU instruction, the SPE APU unavailable exception is taken. 1 Software can execute any of the SPE APU instructions. Embedded floating-point instructions require MSR[SPE] to be set. An attempt to execute an embedded floating-point instruction when MSR[SPE] is 0 results in an SPE APU unavailable interrupt.
39–43	—	Reserved, should be cleared <sup>(1)</sup> .
44	AP	APU available. Book E defines the operation of AP as follows: 0 The processor cannot execute APU instructions. 1 The processor can execute APU instructions.
45	WE	Wait state enable. Allows the core complex to signal a request for power management, according to the states of HID0[DOZE], HID0[NAP], and HID0[SLEEP]. 0 The processor is not in wait state and continues processing. No power management request is signaled to external logic. 1 The processor enters wait state by ceasing to execute instructions and entering low-power mode. Details of how wait state is entered and exited and how the processor behaves in the wait state are implementation-dependent.
46	CE	Critical enable 0 Critical input and watchdog timer interrupts are disabled. 1 Critical input and watchdog timer interrupts are enabled.
47	—	Preserved for Book III ILE

Table 20. MSR field descriptions (continued)

Bits	Name	Description
48	EE	External enable 0 External input, decremter, fixed-interval timer, and performance monitor interrupts are disabled. 1 External input, decremter, fixed-interval timer, and performance monitor interrupts are enabled.
49	PR	User mode (problem state) 0 The processor is in supervisor mode, can execute any instruction, and can access any resource (for example, GPRs, SPRs, and the MSR). 1 The processor is in user mode, cannot execute any privileged instruction, and cannot access any privileged resource. PR also affects memory access control.
50	FP	Floating-point available. 0 The processor cannot execute floating-point instructions, including floating-point loads, stores, and moves. 1 The processor can execute floating-point instructions.
51	ME	Machine check enable. 0 Machine check interrupts are disabled. 1 Machine check interrupts are enabled.
52	FE0	Floating-point exception mode 0. The Book E definition of this bit is shown in <a href="#">Table 21</a> .
53	—	Allocated for implementation-dependent use.
54	DE	Debug interrupt enable 0 Debug interrupts are disabled. 1 Debug interrupts are enabled if DBCR0[IDM] = 1. See the description of the DBSR[UDE] in <a href="#">Section 3.13.2: Debug status register (DBSR)</a> .
55	FE1	Floating-point exception mode 1. The Book E definition of this bit is shown in <a href="#">Table 21</a> .
56	—	Reserved, should be cleared <sup>(1)</sup> .
57	—	Preserved for Book III IP
58	IS	Instruction address space 0 The processor directs all instruction fetches to address space 0 (TS = 0 in the relevant TLB entry). 1 The processor directs all instruction fetches to address space 1 (TS = 1 in the relevant TLB entry).
59	DS	Data address space 0 The processor directs data memory accesses to address space 0 (TS = 0 in the relevant TLB entry). 1 The processor directs data memory accesses to address space 1 (TS = 1 in the relevant TLB entry).
60	—	Reserved, should be cleared <sup>(1)</sup> .

Table 20. MSR field descriptions (continued)

Bits	Name	Description
61	PMM	(Performance monitor APU) Performance monitor mark bit. System software can set PMM when a marked process is running to enable statistics gathering only during the execution of the marked process. PMM and MSR[PR] together define a state that the processor (supervisor or user) and the process (marked or unmarked) may be in at any time. If this state matches an individual state specified in the PMLCax, the state for which monitoring is enabled, counting is enabled.
62–63	—	Preserved for Book III RI and LE, respectively.

1. An MSR bit that is reserved may be altered by return from interrupt instructions.

The floating-point exception mode bits FE0 and FE1 are described in [Table 21](#).

Table 21. Floating-point exception bits—MSR[FE0,FE1]

FE0	FE1	Mode
0	0	Ignore exceptions
0	1	Imprecise nonrecoverable
1	0	Imprecise recoverable
1	1	Precise

## 3.7 Hardware implementation-dependent registers

Each ST Book E processor implements hardware implementation-dependent registers, HID0 and HID1, which contain fields defined either by the EIS or by the implementation. This section provides architectural information about HID registers and describes only those bits that are defined by the EIS.

*Note:* *Not all processors implement HID fields defined by the EIS. Consult the user documentation.*

*An integrated device may not use all HID fields implemented on an embedded core or may define those fields more specifically. Always begin by looking at the core register descriptions in the reference manual for the integrated device.*

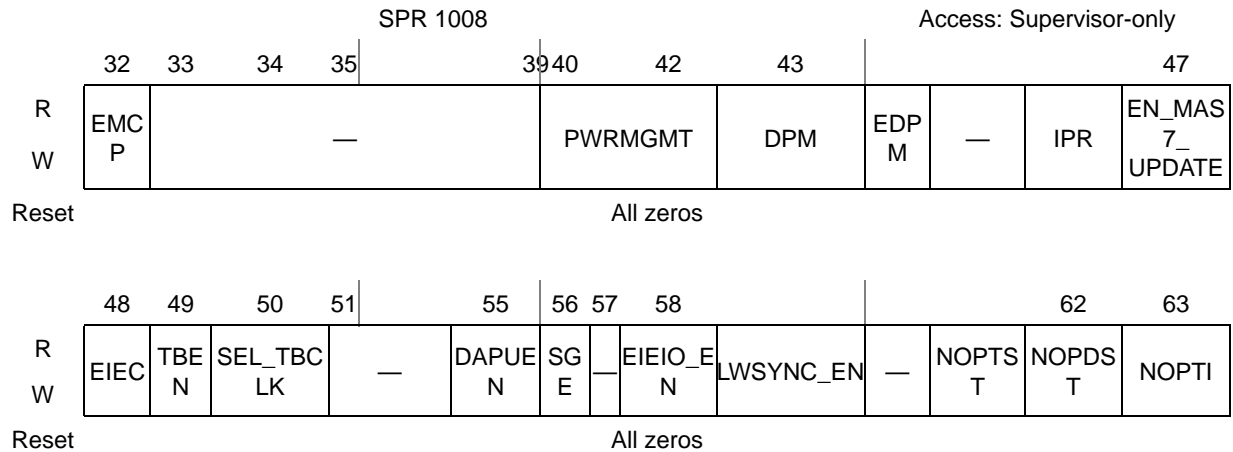
### 3.7.1 Hardware implementation dependent register 0 (HID0)

HID0 is used for configuration and control. Figure below shows the HID0 bits that are defined either generally by the EIS or as part of an EIS-defined APU. Note that not all EIS-compliant device implement all HID0 fields; see the user documentation.

Writing to HID0 typically requires synchronization, as described in [Section 3.18.2: Synchronization requirements for SPRs](#).



Figure 11. Hardware implementation dependent register 0 (HID0)



HID0 fields are described in [Table 22](#).

Table 22. HID0 field descriptions

Bits	Name	Description
32	EMCP	Enable machine check pin. Used to mask machine check exceptions delivered to the core from the machine check input. 0 Machine check exceptions from the machine check signal are disabled. 1 Machine check exceptions from the machine check signal are enabled. If MSR[ME] = 0, asserting the machine signal check causes a checkstop. If MSR[ME] = 1, asserting the machine check signal causes a machine check exception.
33	—	Implementation dependent.
34	SFR	Sixty-four bit results. Determines how the upper 32 bits of 64-bit registers in a 64-bit implementation are computed when the processor is executing in 32-bit mode (MSR[CM] = 0). 0 In 32-bit mode, bits 0–31 of all 64-bit registers are not modified. Explicit 64-bit instructions generate an unimplemented instruction exception when executed. 1 In 32-bit mode, bits 0–31 are written with the same value that is written as when the processor is executing in 64-bit mode (except for the LR and any EAs generated that clear bits 0–31. Explicit 64-bit instructions are allowed to execute and do not generate an unimplemented instruction exception unless they would have when the processor is in 64-bit mode.
35–39	—	Implementation dependent.
40–42	PWRMGMT	Power management control. The semantics of PWRMGMT are implementation dependent.
43	DPM	Dynamic power management. Used to enable power-saving by shutting off functional resources not in use. Setting or clearing DPM should not affect performance. 0 Dynamic power management is disabled. 1 Dynamic power management is enabled.
44	EDPM	Enhanced dynamic power management. Used to enable additional power-saving by shutting off functional resources not in use. Setting EDPM may have adverse effects on performance. 0 Enhanced dynamic power management is disabled. 1 Enhanced dynamic power management is enabled.

Table 22. HID0 field descriptions (continued)

Bits	Name	Description
45	—	Implementation dependent.
46	ICR	Interrupt inputs clear reservation. Controls whether external input and critical input interrupts cause an established reservation to be cleared. 0 External and critical input interrupts do not affect reservation status. 1 External and critical input interrupts, when taken, clear an established reservation.
47	EN_MAS7_UPD ATE	Enable hot-wire update of MAS7 register. Implementations that support this bit do not update MAS7 (upper RPN field) when hardware writes MAS registers via a <b>tlbre</b> , <b>tlbsx</b> , or an interrupt unless this bit is set. This provides a compatibility path for processors that originally offered only 32 bits of physical addressing but have since extended past 32 bits. 0 Hardware updates of MAS7 are disabled. 1 Hardware updates of MAS7 are enabled.
48	EIEC	Enable internal error checking. Used to control whether internal processor errors cause a machine check exception. 0 Internal error reporting is disabled. Internally detected processor errors do not generate a machine check interrupt. 1 Internal error reporting is enabled. Internally detected processor errors generate a machine check interrupt.
49	TBEN	Time base enable. Used to control whether the time base increments. 0 The time base is not enabled and will not increment. 1 The time base is enabled and will increment. The rate at which the time base increments is determined by the value of HID0[SEL_TBCLK].
50	SEL_TBCLK	Select time base clock. Used to select the source of the time base clock. 0 The time base is updated based on a core implementation specific rate. 1 The time base is updated based on an external signal to the core
51–54	—	Implementation dependent.
55	DAPUEN	Debug APU enable. Controls whether the debug APU or enhanced debug APU is enabled. 0 The debug APU is disabled. Debug interrupts use CSRR0 and CSRR1 to save state and the <b>rfdi</b> instruction to return from the debug interrupt. 1 The debug APU is enabled; debug interrupts use DSRR0 and DSRR1 to save state and the <b>rfdi</b> instruction to return from the debug interrupt.
56	SGE	Store gathering enable. Turns on store gathering for non-guarded cache inhibited or write-through stores. Details and characteristics of how stores are gathered is implementation dependent. 0 Store gathering is disabled. 1 Store gathering is enabled.
57	—	Implementation dependent.
58	EIEIO_EN	<b>eiio</b> synchronization enable. Allows <b>mbar</b> instructions to provide the same synchronization semantics as the <b>eiio</b> instruction. 0 Synchronization provided by <b>mbar</b> is performed in the Book E manner. Additional forms of synchronization, if implemented, are determined by the MO value. 1 Synchronization provided by <b>mbar</b> is equivalent to <b>eiio</b> synchronization. The MO field is ignored.


Table 22. HID0 field descriptions (continued)

Bits	Name	Description
59	LWSYNC_EN	Lightweight synchronization enable. Allows <b>msync</b> instructions to provide the same synchronization semantics as the <b>sync</b> instructions from the PowerPC 2.xx architecture. 0 The synchronization provided by the <b>msync</b> instruction is performed in the Book E manner. 1 The synchronization provided by the <b>msync</b> instruction is based on the L field defined in PowerPC 2.xx architecture <b>sync</b> instruction.
60	—	Implementation dependent.
61	NOPTST	No-op cache touch for store instructions. Controls whether data cache touch for store instructions perform no operation. 0 <b>dcbtst</b> , <b>dstst</b> , and <b>dststt</b> and other forms of cache touch for store instructions operate as defined by the EIS and Book E unless disabled by NOPDST or NOPTI. 1 <b>dcbtst</b> , <b>dstst</b> , and <b>dststt</b> and other forms of cache touch for store instructions are treated as no-ops. Cache line touch for store and lock instructions defined in the cache line locking APU operate as defined.
62	NOPDST	No-op <b>dst</b> , <b>dstt</b> , <b>dstst</b> , and <b>dststt</b> instructions. Instructions that start data stream prefetching through the <b>dst</b> instructions produce no-operation. 0 <b>dst</b> , <b>dstt</b> , <b>dstst</b> , and <b>dststt</b> operate as defined by the EIS unless disabled by NOPDST or NOPTI. 1 <b>dst</b> , <b>dstt</b> , <b>dstst</b> , and <b>dststt</b> are treated as no-ops and all current <b>dst</b> prefetch streams are terminated.
63	NOPTI	No-op cache touch instructions. Data and instruction cache touch instructions perform no operations. 0 <b>dcbt</b> , <b>dcbtst</b> , <b>icbt</b> and other forms of cache touch instructions operate as defined by the EIS and Book E unless disabled by NOPDST or NOPTST. 1 <b>dcbt</b> , <b>dcbtst</b> , <b>icbt</b> and other cache touch instruction forms are treated as no-ops. Cache line touch and lock instructions defined in the cache line locking APU operate as defined.

### 3.7.2 Hardware implementation dependent register 1 (HID1)

The EIS defines a HID1 register. HID1 contents are implementation dependent. HID1 is used for bus configuration and control. Writing to HID1 requires synchronization, as described in [Section 3.18.2: Synchronization requirements for SPRs](#).


**Figure 12. Hardware implementation dependent register 1 (HID1)**



**3.7.3 Processor ID register (PIR)**

The processor ID register (PIR), shown below, contains a value that can be used to distinguish the processor from other processors in the system.


**Figure 13. Processor ID register (PIR)**



**3.7.4 Processor version register (PVR)**

The read-only processor version register (PVR), contains a value identifying the version and revision level of the processor. The PVR distinguishes between processors that differ in attributes that may affect software.

**Figure 14. Processor version register (PVR)**



[Table 23](#) describes PVR fields.


**Table 23. PVR field descriptions**

Bits	Name	Description
32–47	Version	A 16-bit number that identifies the version of the processor. Different version numbers indicate major differences between processors, such as which optional facilities and instructions are supported.
48–63	Revision	A 16-bit number that distinguishes between implementations of the version. Different revision numbers indicate minor differences between processors having the same version number, such as clock rate and engineering change level.

### 3.7.5 System version register (SVR)

The system version register (SVR), contains a read-only SoC-dependent value; consult the documentation for the implementation.


**Figure 15. System version register (SVR)**



### 3.8 Timer registers

The time base (TB), decremter (DEC), fixed-interval timer (FIT), and watchdog timer provide timing functions for the system. The relationship of these timer facilities to each other is shown in [Figure 16](#) and is described as follows:

**Figure 16. Relationship of timer facilities to the time base**



- The TB is a long-period counter driven at an implementation-dependent frequency.
- The decremter, updated at the same rate as the TB, provides a way to signal an exception after a specified period unless one of the following occurs:
  - DEC is altered by software in the interim.
  - The TB update frequency changes.
- The DEC is typically used as a general-purpose software timer.
- The time base for the TB and DEC is selected by the time base enable (TBEN) and select time base clock (SEL\_TBCLK) bits in HID0, as follows:
  - If HID0[TBEN] = 1 and HID0[SEL\_TBCLK] = 0, the time base is updated every 8 bus clocks.
  - If HID0[TBEN] = 1 and HID0[SEL\_TBCLK] = 1, the time base is updated by an implementation-specific clock input).
- Software can select one from of four TB bits to signal a fixed-interval interrupt whenever the bit transitions from 0 to 1. It is typically used to trigger periodic system maintenance functions. Bits that may be selected are implementation-dependent.
- The watchdog timer, also a selected TB bit, provides a way to signal a critical exception when the selected bit transitions from 0 to 1. It is typically used for system error recovery. If software does not respond in time to the initial interrupt by clearing the associated status bits in the TSR before the next expiration of the watchdog timer interval, a watchdog timer-generated processor reset may result, if so enabled.


All timer facilities must be initialized during start-up.

### 3.8.1 Timer control register (TCR)

The TCR, provides control information for the on-chip timer of the core complex. The core complex implements two fields not specified in Book E: TCR[WPEXT] and TCR[FPEXT].

The 32-bit timer control register (TCR), controls the decremter (see [Section 3.8.4: Decrementer register](#)).

**Figure 17. Timer control register (TCR)**



[Table 24](#) describes the TCR fields.

**Table 24. TCR field descriptions**

Bits	Name	Description
32–33	WP	Watchdog timer period. When concatenated with WPEXT, specifies one of 64-bit locations of the time base used to signal a watchdog timer exception on a transition from 0 to 1. WPEXT,WP = 0000_00 selects TBU[32] (the msb of the TB) WPEXT,WP = 1111_11 selects TBL[63] (the lsb of the TB)
34–35	WRC	Watchdog timer reset control. When a watchdog reset event occurs, the value programmed into WRC is reflected on <i>core_wrs</i> and into TSR[WRS], but the WRC bits are reset to 00. At this point, software can reprogram WRC. Although WRC can be set by software, it cannot be cleared by software (except by a software-induced reset). Once written to a non-zero value, WRC may no longer be altered by software. 00 No watchdog timer reset will occur. TCR[WRC] resets to 00; it can be set by software, but cannot be cleared by software (except by a software-induced reset). xx Other values: Force processor to be reset on second time-out of watchdog timer. The exact function of any of these settings is implementation-dependent.
36	WIE	Watchdog timer interrupt enable 0 Watchdog timer interrupts disabled 1 Watchdog timer interrupts enabled
37	DIE	Decrementer interrupt enable 0 Decrementer interrupts disabled 1 Decrementer interrupts enabled
38–39	FP	Fixed interval timer period. When concatenated with FPEXT, FP specifies one of 64 bit locations of the time base used to signal a fixed-interval timer exception on a transition from 0 to 1. FPEXT  FP = 0000_00 selects TBU[32] (the msb of the TB) FPEXT  FP = 1111_11 selects TBL[63] (the lsb of the TB)

Table 24. TCR field descriptions (continued)

Bits	Name	Description
40	FIE	Fixed interval interrupt enable 0 Fixed interval interrupts disabled 1 Fixed interval interrupts enabled
41	ARE	Auto-reload enable. Controls whether the value in DECAR is reloaded into the DEC when the DEC value reaches 0000_0001. 0 Auto-reload disabled 1 Auto-reload enabled
42	—	Reserved, should be cleared.
43–46	WPEXT	(EIS) Watchdog timer period extension (see the description for WP)
47–50	FPEXT	(EIS) Fixed-interval timer period extension (see the description for FP)
51–63	—	Reserved, should be cleared.

### 3.8.2 Timer status register (TSR)

As shown below, the 32-bit TSR contains status on timer events and the most recent watchdog timer-initiated processor reset. All TSR bits function as write-1-to-clear.

*Note: Register fields designated as write-1-to-clear are cleared only by writing ones to them. Writing zeros to them has no effect.*

Figure 18. Timer status register (TSR)




Table 25 describes TSR fields.

Table 25. TSR field descriptions

Bits	Name	Description
32	ENW	Enable next watchdog time. When a watchdog timer time-out occurs while WIS = 0 and the next watchdog time-out is enabled (ENW = 1), a watchdog timer exception is generated and logged by setting WIS. This is referred to as a watchdog timer first time out. A watchdog timer interrupt occurs if enabled by TCR[WIE] and MSR[CE]. To avoid another watchdog timer interrupt once MSR[CE] is reenabled (assuming TCR[WIE] is not cleared instead), the interrupt handler must reset TSR[WIS] by executing <b>an mtspr</b> , setting WIS and any other bits to be cleared and a 0 in all other bits. The data written to the TSR is not direct data, but a mask. A 1 causes the bit to be cleared; a 0 has no effect. 0 Action on next watchdog timer time-out is to set TSR[ENW]. 1 Action on next watchdog timer time-out is governed by TSR[WIS].




**Table 25. TSR field descriptions (continued)**

Bits	Name	Description
33	WIS	Watchdog timer interrupt status. See the ENW description for more information about how WIS is used. 0 A watchdog timer event has not occurred. 1 A watchdog timer event occurred. When MSR[CE] = 1 and TCR[WIE] = 1, a watchdog timer interrupt is taken.
34–35	WRS	Watchdog timer reset status. Defined at reset (value = 00). Set to TCR[WRC] when a reset is caused by the watchdog timer. 00 No watchdog timer reset has occurred. xx All other values are implementation-dependent.
36	DIS	Decrementer interrupt status. 0 A decremter event has not occurred. 1 A decremter event occurred. When MSR[EE] = TCR[DIE] = 1, a decremter interrupt is taken.
37	FIS	Fixed-interval timer interrupt status. 0 A fixed-interval timer event has not occurred. 1 A fixed-interval timer event occurred. When MSR[EE] = 1 and TCR[FIE] = 1, a fixed-interval timer interrupt is taken.
38–63	—	Reserved, should be cleared.

### 3.8.3 Time base (TBU and TBL)

The time base (TB), seen below, is composed of two 32-bit registers, the time base upper (TBU) concatenated on the right with the time base lower (TBL). TB provides timing functions for the system. TB is a volatile resource and must be initialized during start-up.

**Figure 19. Time base upper/lower registers (TBU/TBL)**



The TB is interpreted as a 64-bit unsigned integer that is incremented periodically. Each increment adds 1 to the least-significant bit. The frequency at which the integer is updated is implementation-dependent.

TBL increments until its value becomes 0xFFFF\_FFFF ( $2^{32} - 1$ ). At the next increment, its value becomes 0x0000\_0000 and TBU is incremented. This process continues until the TBU value becomes 0xFFFF\_FFFF and value TBL value becomes 0xFFFF\_FFFF (TB is interpreted as 0xFFFF\_FFFF\_FFFF\_FFFF ( $2^{64} - 1$ )). At the next increment, the TBU value becomes 0x0000\_0000 and the TBL value becomes 0x0000\_0000. There is no interrupt (or any other indication) when this occurs.

The period depends on the driving frequency. For example, if TB is driven by 100 MHz divided by 32, the TB period is as follows:

$$T_{TB} = 2^{64} \times \frac{32}{10\text{MHz}} = 5.90 \times 10^{12} \text{ seconds}$$

(approximately 187,000 years)

The TB is implemented such that the following requirements are satisfied:

- Loading a GPR from the TB has no effect on the accuracy of the TB
- Storing a GPR to the TB replaces the value in the TB with the value in the GPR

Book E does not specify a relationship between the frequency at which the TB is updated and other frequencies, such as the CPU clock or bus clock in a Book E system. The TB update frequency is not required to be constant. One of the following is required to ensure that system software can keep time of day and operate interval timers:

- The system provides an (implementation-dependent) interrupt to software whenever the update frequency of the TB changes and a way to determine the current update frequency.
- The update frequency of the TB is under the control of system software.

*Note: Disabling the TB or making reading the time base privileged prevents the TB from being used to implement a covert channel in a secure system.*

*If the operating system initializes the TB on power-on to some reasonable value and the update frequency of the TB is constant, the TB can be used as a source of values that increase at a constant rate, such as for time stamps in trace entries.*

*Even if the update frequency is not constant, values read from the TB are monotonically increasing (except when the TB wraps from  $2^{64} - 1$  to 0). If a trace entry is recorded each time the update frequency changes, the sequence of TB values can be post-processed to become actual time values.*

*Successive readings of the TB may return identical values.*

It is intended that the TB be useful for timing reasonably short sequences of code (a few hundred instructions) and for low-overhead time stamps for tracing.


### 3.8.4 Decrementer register

The 32-bit decrementer (DEC), shown below, is a decrementing counter that is updated at the same rate as the TB. It provides a way to signal a decrementer interrupt after a specified period unless one of the following occurs:

- DEC is altered by software in the interim.
- The TB update frequency changes.

DEC is typically used as a general-purpose software timer. The decrementer auto-reload register is used to automatically reload a programmed value into DEC, as described in [Section 3.8.5: Decrementer auto-reload register \(DECAR\)](#).


**Figure 20. Decrementer register (DEC)**



### 3.8.5 Decrementer auto-reload register (DECAR)

The decrementer auto-reload register is shown in figure below. If the auto-reload function is enabled (TCR[ARE] = 1), the auto-reload value in DECAR is written to DEC when DEC decrements from 0x0000\_0001 to 0x0000\_0000. Note that writing DEC with zeros by using an `mtspr[DEC]` does not automatically generate a decrementer exception.

**Figure 21. Decrementer auto-reload register (DECAR)**



## 3.9 Interrupt registers

[Section 3.9.1: Interrupt registers defined by book E](#), describes registers used for interrupt handling.

### 3.9.1 Interrupt registers defined by book E


This section describes the following register bits and their fields:

- [Section 3.9.1.1: Save/restore register 0 \(SRR0\)](#)
- [Section 3.9.1.2: Save/restore register 1 \(SRR1\)](#)
- [Section 3.9.1.3: Critical save/restore register 0 \(CSRR0\)](#)
- [Section 3.9.1.4: Critical save/restore register 1 \(CSRR1\)](#)
- [Section 3.9.1.5: Data exception address register \(DEAR\)](#)
- [Section 3.9.1.6: Interrupt vector prefix register \(IVPR\)](#)
- [Section 3.9.1.7: Interrupt vector offset registers \(IVORs\)](#)
- [Section 3.9.1.8: Exception syndrome register \(ESR\)](#)

#### 3.9.1.1 Save/restore register 0 (SRR0)

On a noncritical interrupt, SRR0, shown in figure below, holds the address of the instruction where the interrupted process should resume. The instruction is interrupt-specific, although for instruction-caused exceptions, it is typically the address of the instruction that caused the interrupt. When `rfi` executes, instruction execution continues at the address in SRR0.


**Figure 22. Save/restore register 0 (SRR0)**



### 3.9.1.2 Save/restore register 1 (SRR1)

SRR1 is provided to save and restore machine state on noncritical interrupts. When a noncritical interrupt is taken, MSR contents are placed in SRR1. When **rfi** executes, SRR1 contents are placed into MSR. SRR1 bits that correspond to reserved MSR bits are also reserved. These registers are not affected by **rftci** or **rfmci**. Reserved MSR bits may be altered by **rftci**, **rftci**, or **rfmci**.


**Figure 23. Save/restore register 1 (SRR1)**



### 3.9.1.3 Critical save/restore register 0 (CSRR0)

CSRR0, is provided to save and restore machine state on critical interrupts. It is used by critical interrupts like SRR0 is used for standard interrupts: to hold the address of the instruction to which control is passed at the end of the interrupt handler. When **rftci** executes, instruction execution continues at the address in CSRR0.


**Figure 24. Critical save/restore register 0 (CSRR0)**



### 3.9.1.4 Critical save/restore register 1 (CSRR1)

CSRR1, is used to save and restore machine state on critical interrupts. When a critical interrupt is taken, MSR contents are placed into CSRR1. When **rftci** executes, CSRR1 contents are restored into the MSR. CSRR1 bits that correspond to reserved MSR bits are also reserved; reserved MSR bits may be altered.


**Figure 25. Critical save/restore register 1 (CSRR1)**



### 3.9.1.5 Data exception address register (DEAR)

DEAR, is loaded with the effective address of a data access (caused by a load, store, or cache management instruction) that results in an alignment, data TLB miss, or DSI exception.


**Figure 26. Data exception address register (DEAR)**



### 3.9.1.6 Interrupt vector prefix register (IVPR)

IVPR is used with IVORs to determine the vector address. IVPR[32–47] provides the high-order 16 bits of the address of the exception processing routines. The 16-bit vector offsets are concatenated to the right of IVPR[32–47] to form the address of the exception processing routine. IVPR[48–63] are reserved.


**Figure 27. Interrupt vector prefix register (IVPR)**



### 3.9.1.7 Interrupt vector offset registers (IVORs)

IVORs, hold the quad-word index from the base address provided by the IVPR for each interrupt type.

**Figure 28. Interrupt vector offset registers (IVOR)**



SPR numbers corresponding to IVOR16–IVOR31 are reserved. IVOR32–IVOR47 and IVOR60–IVOR63 are reserved. SPR numbers for IVOR32–IVOR63 are allocated for implementation-dependent use. IVOR assignments are shown in [Table 26](#).

Table 26. IVOR assignments

IVOR Number	SPR	Interrupt type
IVOR0	400	Critical input
IVOR1	401	Machine check
IVOR2	402	Data storage
IVOR3	403	Instruction storage
IVOR4	404	External input
IVOR5	405	Alignment
IVOR6	406	Program
IVOR7	407	Floating-point unavailable
IVOR8	408	System call
IVOR9	409	Auxiliary processor unavailable (optional)
IVOR10	410	Decrementer
IVOR11	411	Fixed-interval timer interrupt
IVOR12	412	Watchdog timer interrupt
IVOR13	413	Data TLB error
IVOR14	414	Instruction TLB error
IVOR15	415	Debug
IVOR16– IVOR31	—	Reserved for future architectural use
IVOR36–IVOR63 allocated for implementation dependent use		
IVOR32	528	SPE APU unavailable
IVOR33	529	(Embedded FP APUs) embedded floating-point data exception
IVOR34	530	(Embedded FP APUs) embedded floating-point round exception
IVOR35	531	(Performance monitor APUs) performance monitor

### 3.9.1.8 Exception syndrome register (ESR)

The ESR, provides a syndrome to differentiate between different kinds of exceptions that can generate the same interrupt type. When such an interrupt is generated, bits corresponding to the specific exception that generated the interrupt are set and all other ESR bits are cleared. Other interrupt types do not affect ESR contents. The ESR does not need to be cleared by software. [Table 27](#) shows ESR bit definitions.

EIS storage defines ESR[DLK] and ESR[ILK] to indicate user cache line locking exceptions, ESR[XTE] for precise external transaction errors, and ESR[EPID] external PID load and store exceptions.

The ESR is defined in Book E. Bits architected by EIS storage are defined here.

Figure 29. Exception syndrome register (ESR)

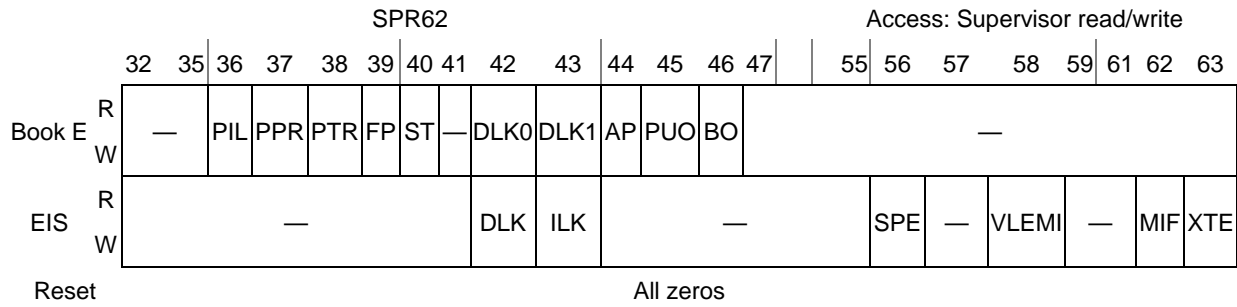


Table 27 describes ESR bit definitions.

Table 27. Exception syndrome register (ESR) definition

Bits	Name	Syndrome	Interrupt types
32–35	—	Reserved, should be cleared. (Defined by Book E as allocated.)	—
36	PIL	Illegal instruction exception	Program
37	PPR	Privileged instruction exception	Program
38	PTR	Trap exception	Program
39	FP	Floating-point operations	Alignment, data storage, data TLB, program
40	ST	Store operation	Alignment, data storage, data TLB error
41	—	Reserved, should be cleared.	—
42	DLK	Defined by cache line locking APU. Instruction cache locking attempt. Set when a DSI occurs because a <b>dcbtls</b> , <b>dcbtstls</b> , or <b>dcblc</b> was executed in user mode (MSR[PR] = 1) while MSR[UCLE] = 0. 0 Default 1 DSI occurred on an attempt to lock line in data cache when MSR[UCLE] = 0.	Data storage
43	ILK	Defined by cache line locking APU. Instruction cache locking attempt. Set when a DSI occurs because an <b>icbtls</b> or <b>icblc</b> was executed in user mode (MSR[PR] = 1) while MSR[UCLE] = 0. 0 Default 1 DSI occurred on an attempt to lock line in instruction cache when MSR[UCLE] = 0.	Data storage
44	APU	Auxiliary processor operation. Defined by Book E.	Alignment, data storage, data TLB, program
45	PUO	Unimplemented operation exception. Defined by Book E.	Program
46	BO	Byte-ordering exception. Defined by Book E and the VLE extension.	Data storage, instruction storage
47	PIE	Imprecise exception. Defined by Book E.	Program
48–55	—	Reserved.	—

Table 27. Exception syndrome register (ESR) definition (continued)

Bits	Name	Syndrome	Interrupt types
56	SPE	Defined by SPE, embedded floating-point APU. SPE/embedded floating-point exception bit 0 Default 1 Any exception caused by an SPE/embedded floating-point instruction occurred.	Data storage, Data TLB error, Alignment, SPE unavailable, Embedded FP unavailable, Embedded FP data, Embedded FP round
57	—	Reserved, should be cleared	
58	VLEMI	Defined by VLE extension. VLEMI indicates that an interrupt was caused by a VLE instruction. VLEMI is set on an exception associated with execution or attempted execution of a VLE instruction. 0 The instruction page associated with the instruction causing the exception does not have the VLE attribute set or the VLE extension is not implemented. 1 The instruction page associated with the instruction causing the exception has the VLE attribute set and the VLE extension is implemented.	Data storage, Data TLB error, Instruction storage, Program, System Call, Alignment, SPE unavailable, Embedded FP unavailable, Embedded FP data, Embedded FP round
59–61	—	Reserved. Defined by Book E as allocated.	—
62	MIF	Defined by the VLE extension. MIF indicates that an interrupt was caused by a misaligned instruction fetch ( $NIA_{62} \neq 0$ ) and the VLE attribute is cleared for the page or the second half of a 32-bit VLE instruction caused an instruction TLB error. 0 Default. 1 $NIA_{62} \neq 0$ and the instruction page associated with NIA does not have the VLE attribute set or the second half of a 32-bit VLE instruction caused an instruction TLB error.	Instruction TLB error, Instruction Storage
63	XTE	External transaction error. An external transaction reported an error but the error was handled precisely by the core. The contents of SRR0 contain the address of the instruction that initiated the transaction. 0 Default. No external transaction error was precisely detected. 1 An external transaction reported an error that was precisely detected.	Instruction storage, Data storage

*Note:* ESR information is incomplete, so system software may need to identify the type of instruction that caused the interrupt and examine the TLB entry and the ESR to fully identify the exception or exceptions. For example, a data storage interrupt may be caused by both a protection violation exception and a byte-ordering exception. System software would have to look beyond ESR[BO], such as the state of MSR[PR] in SRR1 and the TLB entry page protection bits to determine if a protection violation also occurred.

### 3.9.1.9 EIS-defined interrupt registers

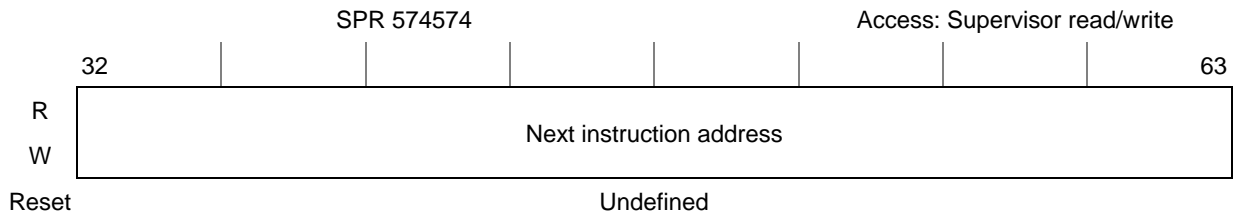
This section describes machine check save/store and syndrome registers.



### 3.9.1.10 Debug save/restore register 0 (DSRR0)

On a debug interrupt, DSRR0, holds the address of the instruction where the interrupted process should resume. The instruction is interrupt specific. See [Section 5.7.16: Debug interrupt](#). When **rfdi** executes, instruction execution continues at the address in DSRR0. DSRR0 and DSRR1 are not affected by **rfi**, **rfci**, or other return from interrupt instructions

**Figure 30. Debug save/restore register 0 (DSRR0)**



### 3.9.1.11 Debug Save/restore register 1 (DSRR1)

DSRR1, is provided to save and restore machine state on debug interrupts. When a debug interrupt is taken, MSR contents are placed into DSRR1. When **rfdi** executes, the contents of DSRR1 are restored into MSR. DSRR1 bits that correspond to reserved MSR bits are also reserved. (See [Section 3.6.1: Machine state register \(MSR\)](#), for more information.) DSRR0 and DSRR1 are not affected by **rfi** or **rfci**. Reserved MSR bits may be altered by **rfi**, **rfci**, or **rfdi**.


**Figure 31. Debug save/restore register 1 (DSRR1)**



### 3.9.1.12 Machine check save/restore register 0 (MCSRR0)

When a machine check interrupt is taken, MCSRR0, is set to the address of the instruction where the interrupted process should resume. The instruction is interrupt-specific, although typically MCSRR0 holds address of the instruction that caused the interrupt. When **rfmci** is executed, instruction execution continues at this address.


**Figure 32. Machine check save/restore register 0 (MCSRR0)**



### 3.9.1.13 Machine check save/restore register 1 (MCSRR1)

MCSRR1 is used to save and restore machine state on machine check interrupts. When a machine check interrupt is taken, MSR contents are placed into MCSRR1. When **rfmci** executes, MCSRR1 contents are restored to MSR. MCSRR1 bits that correspond to reserved MSR bits are also reserved; reserved MSR bits may be altered.


**Figure 33. Machine check save/restore register 1 (MCSRR1)**



### 3.9.1.14 Machine check address register (MCAR/MCARU)

When the core complex takes a machine check interrupt, it updates MCAR, to indicate the address of the data associated with the machine check. Note that if a machine check interrupt is caused by a signal, MCAR contents are not meaningful. Errors that cause MCAR contents to be updated are implementation-dependent. If MCSRR1[MAV] = 1, the address is an effective address; if MAV = 0, the address is a real address.

**Figure 34. Machine check address register (MCAR/MCARU)**



For 32-bit implementations that support physical addresses greater than 32 bits, MCARU provides an alias to the upper address bits that reside in MCAR[0-31].

### 3.9.1.15 Machine check syndrome register (MCSR)

The MCSR, is used to record the cause of the machine check interrupt. In general, machine check syndrome bits correlating to specific hardware error conditions are implementation dependent. Consult the users manual for a complete definition of machine check error syndromes for a specific processor.

**Figure 35. Machine check syndrome register 1 (MCSR)**




Table 28 describes the MCSR fields.

**Table 28. MCSR field descriptions**

Bits	Name	Description
32	MCP	Machine check input to core. Processor cores with a machine check input pin (signal) respond to a signal input by producing an asynchronous machine check. The existence of such a signal and how such a signal is generated is implementation dependent and may be tied to an external pin on the IC package.
33–42	—	Implementation-dependent.
43	NMI	Nonmaskable Interrupt. Set if a non-maskable interrupt (NMI) has been sent to the virtual processor.
44	MAV	MCAR address valid. The address contained in MCAR was updated by the processor and corresponds to the first detected error condition that contained an associated address. Any subsequent machine check errors that have associated addresses are not placed in MCAR unless MAV is 0 when the error is logged. 0 The address in MCAR is not valid. 1 The address in MCAR is valid. Note: Software should read MCAR before clearing MAV. MAV should be cleared before setting MSR[ME].
45	MEA	MCAR effective address. Denotes the type of address in MCAR. MEA has meaning only if MCSR[MAV] is set. 0 The address in MCAR is a physical address. 1 The address in MCAR is an effective address (untranslated).
46–63	—	Implementation-dependent.

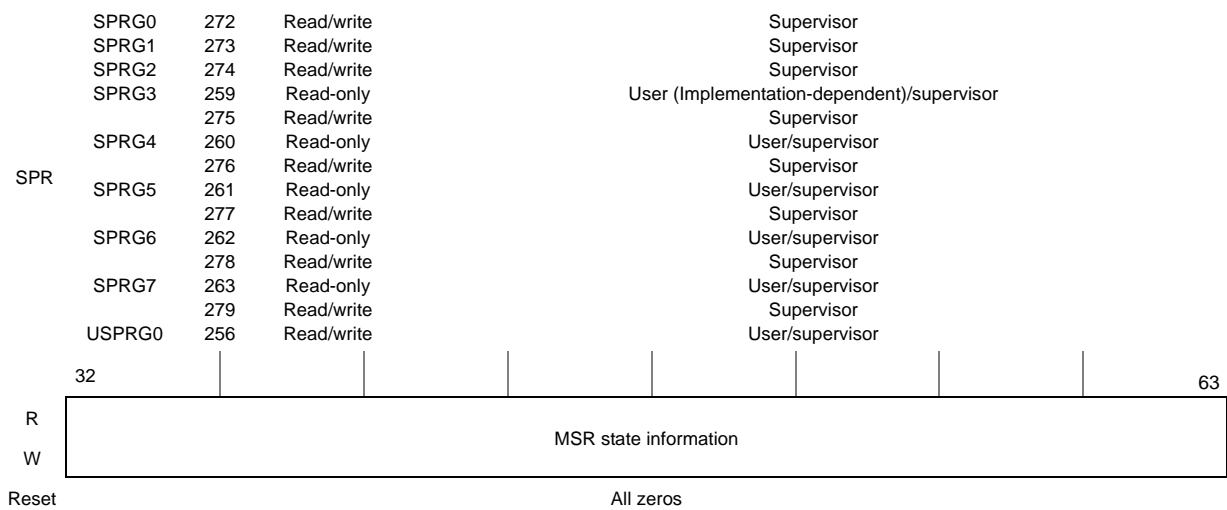
*Note: The machine check interrupt handler should always write what is read back to the MCSR after the error information has been logged. Writing contents that were read from the MCSR back to the MCSR clears only those status bits that were previously read. Failure to clear all MCSR bits causes an asynchronous machine check interrupt when MSR[ME] is set.*

### 3.10 Software use sprs (SPRG0–SPRG7 and USPRG0)

Software-use SPRs (SPRG0–SPRG7 and USPRG0), have no defined functionality. These are shown below:

- SPRG0–SPRG2—can be accessed only in supervisor mode.
- SPRG3—can be written only in supervisor mode. It is readable in supervisor mode, but whether it can be read in user mode is implementation-dependent.
- SPRG4–SPRG7—can be written only in supervisor mode; readable in supervisor or user mode.
- USPRG0—can be accessed in supervisor or user mode.

**Figure 36. Software-use sprs (SPRG0–SPRG7 and USPRG0)**



Software-use SPRs are read into a GPR by using **mf spr** and are written by using **mt spr**.

### 3.11 L1 cache registers

The EIS defines registers that provide control and configuration and status information for the L1 cache implementation.

#### 3.11.1 L1 cache control and status register 0 (L1CSR0)

The L1CSR0, is defined by the EIS. It is used for general control and status of the L1 data cache.

**Figure 37. L1 cache control and status register 0 (L1CSR0)**

SPR 1010 Supervisor read/write

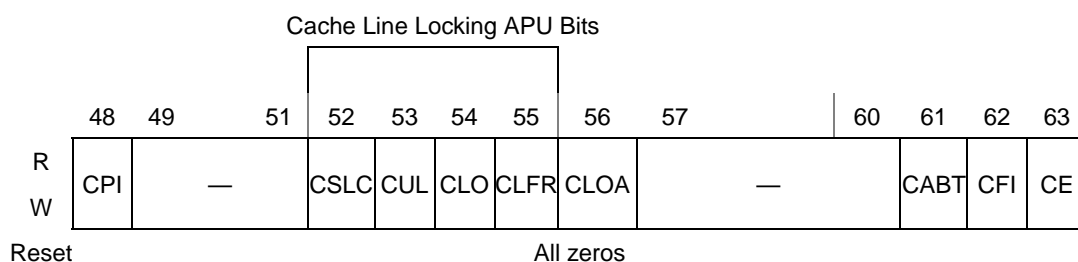
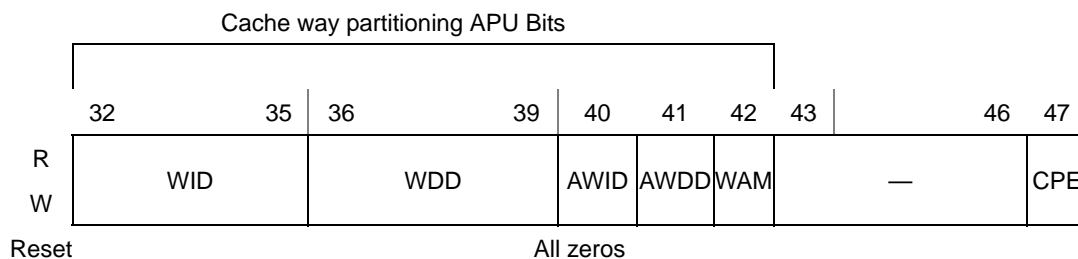


Table 29 describes the L1CSR0 fields.

**Table 29. L1CSR0 field descriptions**

Bits	Name	Description
32–35	WID	Cache way partitioning APU. Way instruction disable. (bit 32 = way 0, bit 33 = way 1, ... bit 35 = way 3). 0 The corresponding way is available for replacement by instruction miss line refills. 1 The corresponding way is not available for replacement by instruction miss line refills.
36–39	WDD	Cache way partitioning APU. Way data disable (bit 36 = way 0, bit 37 = way 1, ... bit 39 = way 3). 0 The corresponding way is available for replacement by data miss line refills. 1 The corresponding way is not available for replacement by data miss line refills
40	AWID	Cache way partitioning APU. Additional ways instruction disable. 0 Additional ways beyond 0–3 are available for replacement by instruction miss line fills. 1 Additional ways beyond 0–3 are not available for replacement by instruction miss line fills.
41	AWDD	Cache way partitioning APU. Additional ways data disable. 0 Additional ways beyond 0–3 are available for replacement by data miss line fills. 1 Additional ways beyond 0–3 are not available for replacement by data miss line fills.
42	WAM	Cache way partitioning APU. Way access mode. 0 All ways are available for access. 1 Only ways partitioned for the specific type of access are used for a fetch or read operation.
43-46	—	Reserved for implementation dependent use.
47	CPE DCPE	[Data] Cache parity enable. 0 Parity checking of the cache disabled 1 Parity checking of the cache enabled

Table 29. L1CSR0 field descriptions (continued)

Bits	Name	Description
48	CPI DCPI	[Data] Cache parity error injection enable. 0 Parity error injection disabled 1 Parity error injection enabled. Note that cache parity must also be enabled (L1CSR0[CPE] = 1) when this bit is set. If DCPE is not set, results are undefined and erratic behavior may occur. It is recommended that an attempt to set this bit when L1CSR0[CPE] = 0 cause the bit not to be set (that is, L1CSR0[CPI] = L1CSR0[CPE] & L1CSR0[CPI]).
49–51	—	Reserved, should be cleared.
52	CSLC DCSLC	[Data]Cache snoop lock clear. Sticky bit set by hardware if a cache line lock was cleared by a snoop operation which caused an invalidation. Note that the lock for that line is cleared whenever the line is invalidated. This bit can be cleared only by software. 0 The cache has not encountered a snoop that invalidated a locked line. 1 The cache has encountered a snoop that invalidated a locked line.
53	CUL DCUL	[Data]Cache unable to lock. Sticky bit set by hardware. This bit can be cleared only by software. 0 Indicates a lock set instruction was effective in the cache 1 Indicates a lock set instruction was not effective in the cache
54	CLO DCLO	[Data]Cache lock overflow. Sticky bit set by hardware. This bit can be cleared only by software. 0 Indicates a lock overflow condition was not encountered in the cache 1 Indicates a lock overflow condition was encountered in the cache
55	CLFC DCLFC	[Data]Cache lock bits flash clear. Clearing occurs regardless of the enable (L1CSR0[CE]) value. 0 Default. 1 Hardware initiates a cache lock bits flash clear operation. Cleared when the operation is complete. During a flash clear operation, writing a 1 causes undefined results; writing a 0 has no effect
56	CLOA DCLOA	[Data]Cache lock overflow allocate. Set by software to allow a lock request to replace a locked line when a lock overflow situation exists. Implementation of this bit is optional. 0 Indicates a lock overflow condition does not replace an existing locked line with the requested line 1 Indicates a lock overflow condition replaces an existing locked line with the requested line
57–60	—	Reserved, should be cleared.
61	CABT DCABT	[Data]Cache operation aborted. 0 No cache operation completed improperly 1 Cache operation did not complete properly
62	CFI DCFI	[Data]Cache flash invalidate. Invalidation occurs regardless of the enable (L1CSR0[CE]) value. 0 No cache invalidate. 1 Cache flash invalidate operation. A cache invalidation operation is initiated by hardware. Once complete, this bit is cleared. During an invalidation operation, writing a 1 causes undefined results; writing a 0 has no effect.
63	CE DCE	[Data]Cache enable. 0 The cache is not enabled. (not accessed or updated) 1 Enables cache operation.

### 3.11.2 L1 cache control and status register 1 (L1CSR1)

L1CSR1, defined as part of the EIS, is used for general control and status of the L1 instruction cache.

**Figure 38. L1 cache control and status register 1 (L1CSR1)**

SPR 10111011

Access: supervisor read/write

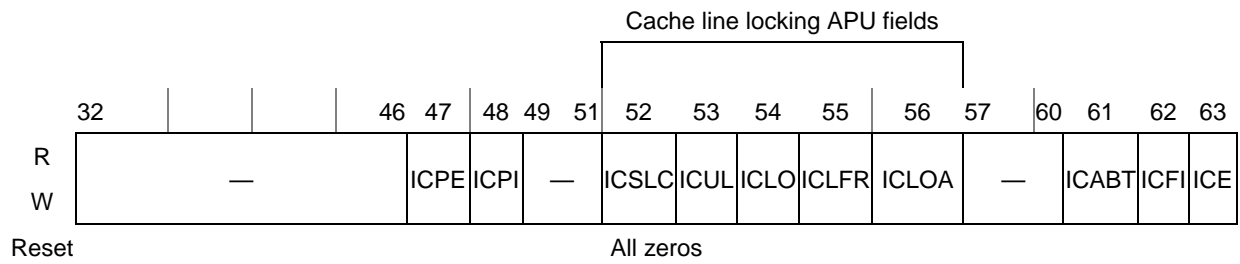


Table 30 describes the L1CSR1 fields.

**Table 30. L1CSR1 field descriptions**

Bits	Name	Description
32–42	—	Reserved, should be cleared.
43–46	—	Reserved for implementation dependent use.
47	ICPE	Instruction cache parity enable. See <a href="#">Section 5.7.2: Machine check interrupt</a> . 0 Parity checking of the cache disabled 1 Parity checking of the cache enabled
48	ICPI	Instruction cache parity error injection enable. 0 Parity error injection disabled 1 Parity error injection enabled. Note that cache parity must also be enabled (L1CSR1[ICPE] = 1) when ICPI is set. If L1CSR0[ICPE] is not set the results are undefined and erratic behavior may occur. It is recommended that an attempt to set this bit when L1CSR0[ICPE] = 0 causes the bit not to be set (that is, L1CSR0[ICPI] = L1CSR0[ICPE] & L1CSR0[ICPI]).
49–51	—	Reserved, should be cleared.
52	ICSLC	Cache line locking APU. Instruction cache snoop lock clear. Sticky bit set by hardware if a cache line lock was cleared by a snoop operation that caused an invalidation. Note that the lock for that line is cleared whenever the line is invalidated. This bit can be cleared only by software. 0 The cache has not encountered a snoop that invalidated a locked line. 1 The cache has encountered a snoop that invalidated a locked line.
53	ICUL	Cache line locking APU. Instruction cache unable to lock. Sticky bit set by hardware. This bit can be cleared only by software. 0 Indicates a lock set instruction was effective in the cache 1 Indicates a lock set instruction was not effective in the cache
54	ICLO DCLO	Cache line locking APU. Instruction cache lock overflow. Sticky bit set by hardware. This bit can be cleared only by software. 0 Indicates a lock overflow condition was not encountered in the cache 1 Indicates a lock overflow condition was encountered in the cache
55	ICLFC	Cache line locking APU. Instruction cache lock bits flash clear. Clearing occurs regardless of the enable (L1CSR1[ICE]) value. 0 Default. 1 Hardware initiates a cache lock bits flash clear operation. This bit is cleared when the operation is complete. During a flash clear operation, writing a 1 causes undefined results; writing a 0 has no effect.


**Table 30. L1CSR1 field descriptions (continued)**

Bits	Name	Description
56	ICLOA	Cache line locking APU. Instruction cache lock overflow no allocate. Set by software to prevent a lock request from replacing a locked line when a lock overflow situation exists. Implementation of this bit is optional. 0 Indicates a lock overflow condition replaces an existing locked line with the requested line 1 Indicates a lock overflow condition does not replace an existing locked line with the requested line
57–60	—	Reserved, should be cleared.
61	ICABT	Instruction cache operation aborted. 0 No cache operation completed improperly 1 Cache operation did not complete properly
62	ICFI	Instruction cache flash invalidate. Invalidation occurs regardless of the enable (L1CSR1[ICE]) value. 0 No cache invalidate. 1 Cache flash invalidate operation. A cache invalidation operation is initiated by hardware. Once complete, this bit is cleared. During an invalidation operation, writing a 1 causes undefined results; writing a 0 has no effect.
63	ICE	Instruction cache enable. 0 The cache is not enabled. (not accessed or updated) 1 Enables cache operation.

**3.11.3 L1 cache configuration register 0 (L1CFG0)**

The L1CFG0 register, shown below, is defined by the EIS to provide configuration information for the primary (L1) data cache of the processor. If a processor implements a unified cache, L1CFG0 applies to the unified cache and L1CFG1 is not implemented.

**Figure 39. L1 cache configuration register 0 (L1CFG0)**



**Table 31. L1CFG0 field descriptions**

Bits	Name	Description
32–33	CARCH	Cache architecture 00 Harvard 01 Unified
34	CWPA	Cache way partitioning APU available. 0 Unavailable 1 Available




**Table 31. L1CFG0 field descriptions (continued)**

Bits	Name	Description
35	CFAHA	Cache flush all by hardware available 0 Unavailable 1 Available
36	CFISWA	Direct cache flush APU available. (Cache flush by set and way available.) 0 Unavailable 1 Available
37–38	—	Reserved, should be cleared.
39–40	CBSIZE	Cache line size 00 32 bytes 01 64 bytes 10 128 bytes 11 Reserved
41–42	CREPL	Cache replacement policy 00 True LRU 01 Pseudo LRU 1x Reserved
43	CLA	Cache line locking APU available 0 Unavailable 1 Available
44	CPA	Cache parity available 0 Unavailable 1 Available
45–52	CNWAY	Cache number of ways minus 1.
53–63	CSIZE	Cache size in Kbytes.

**3.11.4 L1 cache configuration register 1 (L1CFG1)**

The L1CFG1 register, provides configuration information for the L1 instruction cache. If a processor implements a unified cache, L1CFG0 applies to the unified cache and L1CFG1 is not implemented.

**Figure 40. L1 cache configuration register 1 (L1CFG1)**



**Table 32. L1CFG1 field descriptions**

Bits	Name	Description
32–38	—	Reserved, should be cleared.
39–40	ICBSIZ	Instruction cache block size 00 32 bytes 01 64 bytes 10 128 bytes 11 Reserved
41–42	ICREPL	Cache replacement policy 00 True LRU 01 Pseudo LRU 1x Reserved
43	ICLA	Cache line locking APU available 0 Unavailable 1 Available
44	ICPA	Cache parity available 0 Unavailable 1 Available
45–52	ICNWAY	Cache number of ways minus 1.
53–63	ICSIZE	Cache size in Kbytes.

### 3.11.5 L1 flush and invalidate control register 0 (L1FINV0)

The direct cache flush APU defines the L1 flush and invalidate control register 0 (L1FINV0), shown in figure below. The direct cache flush APU allows the programmer to flush and/or invalidate the cache by specifying the cache set and cache way. The direct cache flush APU available bit, L1CFG0[CFISWA], is set for implementations that contain the direct cache flush APU.

To address a specific physical block of the cache, the L1FINV0 is written with the cache set (L1FINV0[CSET]) and cache way (L1FINV0[CWAY]) of the line that is to be flushed. No tag match in the cache is required.

Only the L1 data cache (or unified cache) is manipulated by the direct cache flush APU. The L1 instruction cache or any other caches in the cache hierarchy are not explicitly targeted by this APU. See [Section 9.2: Direct cache flush APU](#).

**Figure 41. L1 flush and invalidate control register 0 (L1FINV0)**




Table 33. L1FINV0 fields—L1 direct cache flush

Bits	Name	Descriptions
0–31	—	Reserved, should be cleared.
32–39	CWAY	Cache way. Specifies the cache way to be selected.
40–41	—	Reserved, should be cleared.
42–58	CSET	Cache set. Specifies the cache set to be selected.
59–61	—	Reserved, should be cleared.
62–63	CCMD	Cache flush command. 00 Implementation dependent. If implemented, the action performed on the line should be synonymous with a <b>dcbi</b> instruction that references the same line. 01 The line specified by CWAY and CSET is flushed if it is modified and valid. It is implementation dependent whether it remains in the cache, or is invalidated. For an implementation, the action performed on the line should be synonymous with a <b>dcbst</b> instruction that references the same line. 01 The line specified by CWAY and CSET is flushed if it is modified and valid. It is then invalidated. For an implementation, the action performed on the line should be synonymous with a <b>dcbf</b> instruction that references that line. 11 Reserved for future use.

## 3.12 MMU registers

This section describes the following MMU registers and their fields:

- Process ID registers (PID0–PID2)
- MMU control and status register 0 (MMUCSR0)
- MMU configuration register (MMUCFG)
- TLB configuration registers (TLB<sub>n</sub>CFG)
- MMU assist registers (MAS0–MAS7)

### 3.12.1 Process ID registers (PID0–PID<sub>n</sub>)

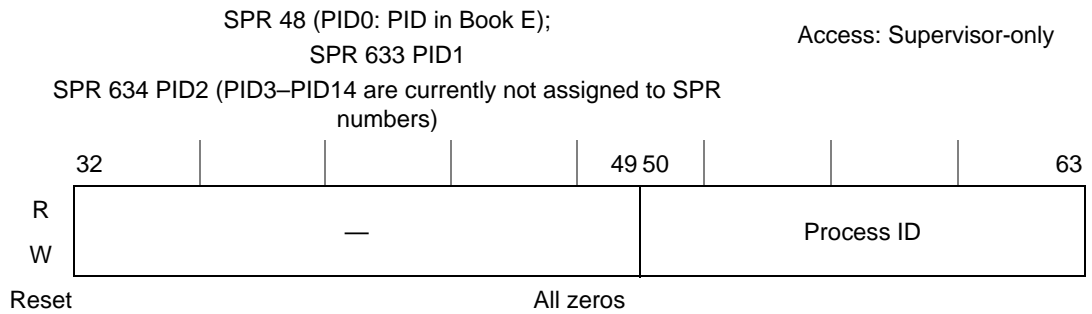
The Book E architecture specifies that a process ID (PID) value be associated with each effective address (instruction or data) generated by the processor.

System software uses PIDs to identify TLB entries that the processor uses to translate addresses for loads, stores, and instruction fetches. PID contents are compared to the TID field in TLB entries as part of selecting appropriate TLB entries for address translation. PID values are used to construct virtual addresses for accessing memory. Note that individual processors may not implement all 14 bits of the process ID field.

Book E defines one PID register that holds the PID value for the current process. ST devices may implement from 1 to 15 PID registers. The number of PIDs implemented is indicated by the value of MMUCFG[NPIDS]. Consult the user documentation for the implementation to determine if other PID registers are implemented.

The suggested PID usage is for PID0 to denote private mappings for a process and for other PIDs to handle mappings that may be common to multiple processes. This method allows for processes sharing address space to also share TLB entries if the shared space is mapped at the same virtual address in each process.


**Figure 42. Process ID registers (PID0–PID2)**



**3.12.2 MMU control and status register 0 (MMUCSR0)**

The MMUCSR0 register is used for general control of the L1 and L2 MMUs.

**Figure 43. MMU control and status register 0 (MMUCSR0)**




**Table 34. MMUCSR0 field descriptions**

Bits	Name	Description
32–60	—	Reserved, should be cleared.
61	L2TLB0_FI TLB0_FI	TLB0 flash invalidate (write 1 to invalidate) 0 No flash invalidate. Writing a 0 to this bit during an invalidation operation is ignored. 1 TLB0 invalidation operation. Hardware initiates a TLB0 invalidation operation. When this operation is complete, this bit is cleared. Writing a 1 during an invalidation operation causes an undefined operation. If the TLB array supports IPROT, entries that have IPROT set are not invalidated.
62	L2TLB1_FI TLB1_FI	TLB1 flash invalidate (write 1 to invalidate) 0 No flash invalidate. Writing a 0 to this bit during an invalidation operation is ignored. 1 TLB1 invalidation operation. Hardware initiates a TLB1 invalidation operation. When this operation is complete, this bit is cleared. Writing a 1 during an invalidation operation causes an undefined operation. This invalidation typically takes 1 cycle.
63	—	Reserved, should be cleared.

### 3.12.3 MMU configuration register (MMUCFG)

MMUCFG, shown below, gives configuration information about the implementation’s MMU.

**Figure 44. MMU configuration register 1 (MMUCFG)**




**Table 35. MMUCFG field descriptions**

Bits	Name	Description
32–48	—	Reserved, should be cleared.
49–52	NPIDS	Number of PID registers, a 4-bit field that indicates the number of PID registers provided by the processor.
53–57	PIDSIZE	PID register size. The PIDSIZE value is one fewer than the number of bits in each PID register implemented. The processor implements only the least significant PIDSIZE+1 bits in the PID registers.
58–59	—	Reserved, should be cleared.
60–61	NTLBS	Number of TLBs. The value of NTLBS is one less than the number of software-accessible TLB structures that are implemented by the processor. NTLBS is set to one less than the number of TLB structures so that its value matches the maximum value of MAS0[TLBSEL].) 00 1 TLB 01 2 TLBs 10 3 TLBs 11 4 TLBs
62–63	MAVN	MMU architecture version number. Indicates the version number of the architecture of the MMU implemented by the processor. 00 Version 1.0 01 Reserved 10 Reserved 11 Reserved

### 3.12.4 TLB configuration registers (TLB $n$ CFG)

TLB $n$ CFG registers, shown below, provide information about each specific TLB that is visible to the programming model. TLB0CFG corresponds to TLB0, TLB1CFG corresponds to TLB1, etc.

**Figure 45. TLB configuration register  $n$  (TLB0CFG)**



**Table 36. TLB $n$ CFG field descriptions**

Bits	Name	Description
32–39	ASSOC	Associativity of TLB $n$ . Number of ways of associativity of TLB array. 0000_0000 Fully associative (A value equal to NENTRY also indicates fully associative.) 0000_0001 1-way set associative 0000_0002 2-way set associative ...
40–43	MINSIZE	Minimum page size of TLB $n$ 0001 Indicates smallest page size is 4 Kbytes 0002 Indicates smallest page size is 8 Kbytes ...
44–47	MAXSIZE	Maximum page size of TLB $n$ 0001 Indicates maximum page size is 4 Kbytes 0002 Indicates maximum page size is 8 Kbytes ...
48	IPROT	Invalidate protect capability of TLB $n$ array. 0 Indicates invalidate protection capability not supported. 1 Indicates invalidate protection capability supported.
49	AVAIL	Page size availability of TLB $n$ array. 0 Fixed selectable page size from MINSIZE to MAXSIZE (all TLB entries are the same size). 1 Variable page size from MINSIZE to MAXSIZE (each TLB entry can be sized separately).
50–51	—	Reserved, should be cleared.
52–63	NENTRY	Number of entries in TLB $n$

### 3.12.5 MMU assist registers (MAS0–MAS7)

MMU assist registers are defined by the EIS and used by the MMU to manage pages and TLBs. Note that some fields in these registers are redefined by implementations.

#### 3.12.5.1 MAS register 0 (MAS0)

MAS0, is used for MMU read/write and replacement control.

Figure 46. MAS register 0 (MAS0)





Table 37. MAS0 field descriptions

Bits	Name	Comments or function when set
32–33	—	Reserved, should be cleared.
34–35	TLBSEL	Selects TLB for access. 00 TLB0 01 TLB1 10 TLB2 11 TLB3
36–47	ESEL	Entry select. Identifies an entry in the selected array to be used for <b>tlbwe</b> and <b>tlbre</b> . Valid values for ESEL are from 0 to TLBnCFG[ASSOC] - 1. That is, ESEL selects the way from a set of entries determined by MAS3[EPN]. For fully associative TLB arrays, ESEL ranges from 0 to TLBnCFG[NENTRY] - 1. ESEL is also updated on TLB error exceptions (misses) and <b>tlbsx</b> hit and miss cases.
48–51	—	Reserved, should be cleared.
52–63	NV	Next victim. For those TLBs that support the NV field, provides a hint to software to identify the next victim to be targeted for a TLB miss replacement operation. If the TLB selected by MAS0[TLBSEL] does not support NV, this field is undefined. The computation of NV is implementation-dependent. NV is updated on TLB error exceptions (misses), <b>tlbsx</b> hit and miss cases, as shown in <a href="#">Table 194</a> , and on execution of <b>tlbre</b> if the accessed TLB array supports NV. If NV is updated by a supported TLB array, NV always presents a value that can be used in MAS0[ESEL].

#### 3.12.5.2 MAS register 1 (MAS1)

Below is the format of MAS1.

**Figure 47. MAS register 1 (MAS1) format**



**Table 38. MAS1 field descriptions—descriptor context and configuration control**

Bits	Name	Descriptions
32	V	TLB valid bit. 0 This TLB entry is invalid. 1 This TLB entry is valid.
33	IPROT	Invalidate protect. Set to protect this TLB entry from invalidate operations due the execution of <b>tlbivax</b> , broadcast invalidations from another processor, or flash invalidations. Note that not all TLB arrays are necessarily protected from invalidation with IPROT. Arrays that support invalidate protection are denoted as such in the TLB configuration registers. 0 Entry is not protected from invalidation. 1 Entry is protected from invalidation.
34–35	—	Reserved, should be cleared.
36–47	TID	Translation identity. During translation, TID is compared with the current process IDs (PIDs) to select a TLB entry. A TID value of 0 defines an entry as global and matches with all process IDs.
48–50	—	Reserved, should be cleared.
51	TS	Translation space. During translation, TS is compared with AS (MSR[IS] or MSR[DS], depending on the type of access) to select a TLB entry.
52–55	TSIZE	Translation size. Defines the page size of the TLB entry. For TLB arrays that contain fixed-size TLB entries, TSIZE is ignored. For variable page-size TLB arrays, the page size is $4^{TSIZE}$ Kbytes. TSIZE must be between TLBnCFG[MINSIZE] and TLBnCFG[MINSIZE]. Note that the EIS standard supports all 16 page sizes defined in Book E.  0001 4 Kbyte      0111 16 Mbyte 0010 16 Kbyte    1000 64 Mbyte 0011 64 Kbyte    1001 256 Mbyte 0100 256 Kbyte   1010 1 Gbyte 0101 1 Mbyte     1011 4 Gbyte 0110 4 Mbyte
56–63	—	Reserved, should be cleared.

**3.12.5.3 MAS register 2 (MAS2)**

MAS2, contains fields for specifying the effective page address and the storage attributes for a TLB entry.



Figure 48. MAS register 2 (MAS2)




Table 39. MAS2 field descriptions—EPN and page attributes

Bits	Name	Description
32–51	EPN	Effective page number. Depending on page size, only the bits associated with a page boundary are valid. Bits that represent offsets within a page are ignored and should be zero. EPN[0–31] are accessible only in 64-bit implementations as the upper 32 bits of the logical address of the page.
52–55	—	Reserved, should be cleared.
56–57	ACM X0	Alternate coherency mode. Allows an implementation to employ multiple coherency methods. If the M attribute (memory coherence required) is not set for a page (M=0), the page has no coherency associated with it and ACM is ignored. If the M attribute is set for a page (M=1), ACM determines the coherency domain (or protocol) used. ACM values are implementation dependent. <b>Note:</b> Some previous implementations may have a storage bit in the bit 57 position labeled as X0.
58	VLE X1	VLE mode. Identifies pages which contain instructions from the VLE instruction set. The VLE attribute is only implemented if the processor supports the VLE extension. Setting the VLE attribute to 1 and setting the E attribute to 1 is considered a programming error and an attempt to fetch instructions from a page so marked produces an instruction storage interrupt byte ordering exception and sets ESR[BO]. 0 Instructions fetched from the page are decoded and executed as PowerPC (and associated EIS APUs) instructions. 1 Instructions fetched from the page are decoded and executed as VLE (and associated EIS APUs) instructions. Implementation-dependent page attribute. <b>Note:</b> Some implementations have a bit in this position labeled as X1. Software should not use the presence of this bit (the ability to set to 1 and read a 1) to determine if the implementation supports the VLE extension.
59	W	Write-through 0 This page is considered write-back with respect to the caches in the system. 1 All stores performed to this page are written through the caches to main memory.
60	I	Caching-inhibited 0 Accesses to this page are considered cacheable. 1 The page is considered caching-inhibited. All loads and stores to the page bypass the caches and are performed directly to main memory. A read or write to a caching-inhibited page affects only the memory element specified by the operation.
61	M	Memory coherence required 0 Memory coherence is not required. 1 Memory coherence is required. This allows loads and stores to this page to be coherent with loads and stores from other processors (and devices) in the system, assuming all such devices are participating in the coherence protocol.


**Table 39. MAS2 field descriptions—EPN and page attributes (continued)**

Bits	Name	Description
62	G	Guarded 0 Accesses to this page are not guarded and can be performed before it is known if they are required by the sequential execution model. 1 Loads and stores to this page are performed without speculation (that is, they are known to be required).
63	E	Endianness. Determines endianness for the corresponding page. Little-endian operation is true little endian, which differs from the modified little-endian byte-ordering model optionally available in previous devices that implement the PowerPC architecture. 0 The page is accessed in big-endian byte order. 1 The page is accessed in true little-endian byte order.

**3.12.5.4 MAS register 3 (MAS3)**

MAS3 contains fields for specifying the real page address and the permission attributes for a TLB entry.

**Figure 49. MAS register 3 (MAS3)**




**Table 40. MAS3 field descriptions—RPN and access control**

Bits	Name	Description
32–51	RPN[32–51]	Real page number bits 32–51. Depending on page size, only the bits associated with a page boundary are valid. Bits that represent offsets within a page are ignored and should be zero. If the physical address space exceeds 32 bits, RPN[0–31] are accessed through MAS7.
52–53	—	Reserved, should be cleared.
54–57	U0–U3	User bits. Associated with a TLB entry and used by system software. For example, these bits may be used to hold information useful to a page scanning algorithm or be used to mark more abstract page attributes.
58–63	UX, SX UW, SW UR, SR	Permission bits (UX, SX, UW, SW, UR, SR). User and supervisor read, write, and execute permission bits. Effects of the permission bits are defined by Book E.

**3.12.5.5 MAS register 4 (MAS4)**

MAS4, contains fields for specifying default information to be pre-loaded on certain MMU related exceptions.

Figure 50. MAS register 4 (MAS4)



The MAS4 fields are described in [Table 41](#).


Table 41. MAS4 field descriptions—hardware replacement assist configuration

Bits	Name	Description
32–33	—	Reserved, should be cleared.
34–35	TLBSELD	TLBSEL default value. Specifies the default value loaded in MAS0[TLBSEL] on a TLB miss exception.
36–43	—	Reserved, should be cleared.
44–47	TIDSELD	TID default selection value. Specifies which of the current PID registers should be used to load MAS1[TID] on a TLB miss exception. PID registers are addressed as follows: 0000 = PID0 (PID) 0001 = PID1 ... 1110 = PID14 A value that references a non-implemented PID register causes a value of 0 to be placed in MAS1[TID]. See the implementations documentation for a list of supported PIDs.
48–51	—	Reserved, should be cleared.
52–55	TSIZED	Default TSIZE value. Specifies the default value loaded into MAS1[TSIZE] on a TLB miss exception.
56–57	ACMD	Default ACM value. Specifies the default value loaded into MAS2[ACM] on a TLB miss exception.
58	VLED	Default VLE value. Specifies the default value loaded into MAS2[VLE] on a TLB miss exception.
59	WD	Default W value. Specifies the default value loaded into MAS2[W] on a TLB miss exception.
60	ID	Default I value. Specifies the default value loaded into MAS2[I] on a TLB miss exception.
61	MD	Default M value. Specifies the default value loaded into MAS2[M] on a TLB miss exception.
62	GD	Default G value. Specifies the default value loaded into MAS2[G] on a TLB miss exception.
63	ED	Default E value. Specifies the default value loaded into MAS2[E] on a TLB miss exception.

**3.12.5.6 MAS register 5 (MAS5)**

The optional MAS5 register, contains fields for specifying PID values to be used when searching TLB entries with the **tlbsx** instruction.

**Figure 51. MAS register 5 (MAS5)**




**Table 42. MAS5 field descriptions—extended search pIDs**

Bits	Name	Description
32–33	—	Reserved, should be cleared.
34–47	SPID2	Search PID2. Specifies the PID2 value used when searching the TLB during execution of <b>tlbsx</b> . This field is optional and if implemented is valid for only the number of bits implemented for PID registers.
48–49	—	Reserved, should be cleared.
50–63	SPID3	Search PID3. Specifies the PID3 value used when searching the TLB during execution of <b>tlbsx</b> . This field is optional and if implemented is valid for only the number of bits implemented for PID registers.

**3.12.5.7 MAS register 6 (MAS6)**

MAS6, contains fields for specifying PID and AS values to be used when searching TLB entries with the **tlbsx** instruction.

**Figure 52. MAS register 6 (MAS6)**



**Table 43. MAS 6 field descriptions—search pids and search AS**

Bits	Name	Description
32–33	—	Reserved, should be cleared.
34–47	SPID0	Search PID0. Specifies the value of PID0 used when searching the TLB during execution of <b>tlbsx</b> . SPID0 is valid for only the number of bits implemented for PID registers.
48	—	Reserved, should be cleared.


**Table 43. MAS 6 field descriptions—search pids and search AS (continued)**

Bits	Name	Description
49–62	SPID1	Search PID1. Specifies the value of PID1 used when searching the TLB during execution of <b>tlbsx</b> . SPID1 is optional, and if implemented is valid for only the number of bits implemented for PID registers.
63	SAS	Address space value for searches. Specifies the AS value used when executing <b>tlbsx</b> to search the TLB.

**3.12.5.8 MAS register 7 (MAS7)**

MAS7, contains the high-order address bits of the RPN only for implementations that support more than 32 bits of physical address.

**Figure 53. MAS register 7 (MAS7)**



**Table 44. MAS 7 field descriptions—high order RPN**

Bits	Name	Description
32–63	RPN[0–31]	Real page number (bits 0–31). RPN[32–63] are accessed through MAS3.

**3.13 Debug registers**

This section describes debug-related registers that are accessible to software running on the processor. These registers are intended for use by special debug tools and debug software, and not by general application or operating system code.

### 3.13.1 Debug control registers (DBCR0–DBCR3)

The debug control registers are used to enable debug events, reset the processor, control timer operation during debug events, and set the debug mode of the processor.

#### 3.13.1.1 Debug control register 0 (DBCR0)

Below is the DBCR0.

Figure 54. Debug control register 0 (DBCR0)

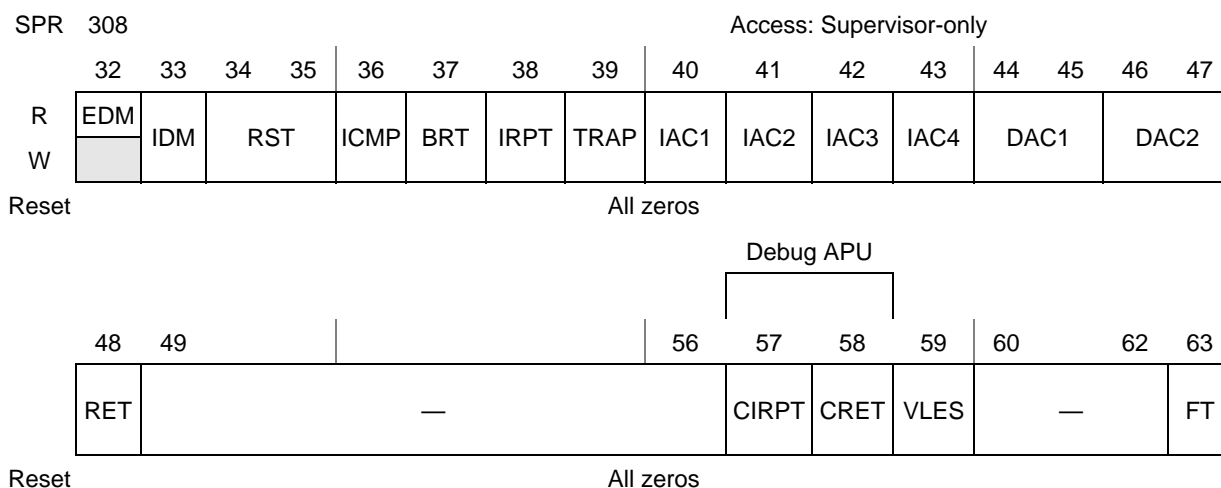


Table 45. DBCR0 field descriptions

Bits	Name	Description
32	EDM	External debug mode. Indicates whether the processor is in external debug mode. 0 The processor is not in external debug mode. 1 The processor is in external debug mode. In some implementations, if EDM = 1, some debug registers are locked and cannot be accessed. Refer to the implementation documentation for any additional implementation-specific behavior.
33	IDM	Internal debug mode. 0 Debug interrupts are disabled. No debug interrupts are taken and debug events are not logged. 1 If MSR[DE] = 1, the occurrence of a debug event or the recording of an earlier debug event in the DBSR when MSR[DE] = 0 or DBCR0[IDM] = 0 causes a debug interrupt.  Programming note: Software must clear debug event status in the DBSR in the debug interrupt handler when a debug interrupt is taken before re-enabling interrupts through MSR[DE]. Otherwise, redundant debug interrupts are taken for the same debug event.
34–35	RST	Reset. Book E defines RST such that 00 is always no action and all other settings are implementation 0x Default (No action) 1x A hard reset is performed on the processor.

Table 45. DBCR0 field descriptions (continued)

Bits	Name	Description
36	ICMP	Instruction completion debug event enable 0 ICMP debug events are disabled. 1 ICMP debug events are enabled. Note: Instruction completion does not cause an ICMP debug event if MSR[DE]=0.
37	BRT	Branch taken debug event enable 0 BRT debug events are disabled. 1 BRT debug events are enabled. Note: Taken branches do not cause a BRT debug event if MSR[DE]=0.
38	IRPT	Interrupt taken debug event enable. 0 IRPT debug events are disabled. 1 IRPT debug events are enabled
39	TRAP	Trap debug event enable 0 TRAP debug events cannot occur. 1 TRAP debug events can occur.
40	IAC1	Instruction address compare 1 debug event enable 0 IAC1 debug events cannot occur. 1 IAC1 debug events can occur.
41	IAC2	Instruction address compare 2 debug event enable. 0 IAC2 debug events cannot occur. 1 IAC2 debug events can occur.
42	IAC3	Defined by Book E as instruction address compare 3 debug event enable 0 IAC3 debug events cannot occur. 1 IAC3 debug events can occur.
43	IAC4	Defined by Book E as instruction address compare 4 debug event enable 0 IAC4 debug events cannot occur. 1 IAC4 debug events can occur.
44–45	DAC1	Data address compare 1 debug event enable 00 DAC1 debug events cannot occur. 01 DAC1 debug events can occur only if a store-type data storage access. 10 DAC1 debug events can occur only if a load-type data storage access. 11 DAC1 debug events can occur on any data storage access.
46–47	DAC2	Data address compare 2 debug event enable 00 DAC2 debug events cannot occur. 01 DAC2 debug events can occur only if a store-type data storage access. 10 DAC2 debug events can occur only if a load-type data storage access. 11 DAC2 debug events can occur on any data storage access.
48	RET	Return debug event enable 0 RET debug events cannot occur. 1 RET debug events can occur. <b>Note:</b> An <b>rfci</b> does not cause an RET debug event if MSR[DE] = 0 at the time that <b>rfci</b> executes.
49–56	—	Reserved, should be cleared.


**Table 45. DBCR0 field descriptions (continued)**

Bits	Name	Description
57	CIRPT	Debug APU, Critical interrupt taken debug event. A critical interrupt taken debug event occurs when DBCR0[CIRPT] = 1 and a critical interrupt (any interrupt that uses the critical class, that is, uses CSRR0 and CSRR1) occurs. 0 Critical interrupt taken debug events are disabled. 1 Critical interrupt taken debug events are enabled.
58	CRET	Debug APU. Critical interrupt return debug event. A critical interrupt return debug event occurs when DBCR0[CRET] = 1 and a return from critical interrupt (an <b>rfci</b> instruction is executed) occurs. 0 Critical interrupt return debug events are disabled. 1 Critical interrupt return debug events are enabled.
59	VLES	VLE status. (VLE APU). Undefined for IRPT, CIRPT, DEVT[1,2], DCNT[1,2], and UDE events. 0 CRET debug events are disabled. 1 An ICMP, BRT, TRAP, RET, CRET, IAC, or DAC debug event occurred on a VLE instruction.
60–62	—	Reserved
63	FT	Freeze timers on debug event 0 Enable clocking of timers. 1 Disable clocking of timers if any DBSR bit is set (except MRR).

**3.13.1.2 Debug control register 1 (DBCR1)**

DBCR1 is shown below.

**Figure 55. Debug control register 1 (DBCR1)**



[Table 46](#) provides bit definitions for the DBCR1.

**Table 46. DBCR1 field descriptions**

Bits	Name	Description
32–33	IAC1US	Instruction address compare 1 user/supervisor mode 00 IAC1 debug events can occur. 01 Reserved 10 IAC1 debug events can occur only if MSR[PR]=0. 11 IAC1 debug events can occur only if MSR[PR]=1.



Table 46. DBCR1 field descriptions (continued)

Bits	Name	Description
34–35	IAC1ER	Instruction address compare 1 effective/real mode 00 IAC1 debug events are based on effective addresses. 01 IAC1 debug events are based on real addresses. 10 IAC1 debug events are based on effective addresses and can occur only if MSR[IS]=0. 11 IAC1 debug events are based on effective addresses and can occur only if MSR[IS]=1.
36–37	IAC2US	Instruction address compare 2 user/supervisor mode 00 IAC2 debug events can occur. 01 Reserved 10 IAC2 debug events can occur only if MSR[PR]=0. 11 IAC2 debug events can occur only if MSR[PR]=1.
38–39	IAC2ER	Instruction address compare 2 effective/real mode 00 IAC2 debug events are based on effective addresses. 01 IAC2 debug events are based on real addresses. 10 IAC2 debug events are based on effective addresses and can occur only if MSR[IS]=0. 11 IAC2 debug events are based on effective addresses and can occur only if MSR[IS]=1.
40–41	IAC12M	Instruction address compare 1/2 mode 00 Exact address compare. IAC1 debug events can occur only if the instruction fetch address equals the value in IAC1. IAC2 debug events can occur only if the instruction fetch address equals the value in IAC2. 01 Address bit match. IAC1 and IAC2 debug events can occur only if the instruction fetch address, ANDed with the contents of IAC2, equals the value in IAC1, also ANDed with the contents of IAC2. If IAC1US≠IAC2US or IAC1ER≠IAC2ER, results are boundedly undefined. 10 Inclusive address range compare. IAC1 and IAC2 debug events can occur only if the instruction fetch address lies between the values specified in IAC1 and IAC2. If IAC1US≠IAC2US or IAC1ER≠IAC2ER, results are boundedly undefined. 11 Exclusive address range compare. IAC1 and IAC2 debug events can occur only if the instruction fetch address lies between the values specified in IAC1 and IAC2. If IAC1US≠IAC2US or IAC1ER≠IAC2ER, results are boundedly undefined.
42–47	—	Reserved, should be cleared.
48–49	IAC3US	Instruction address compare 3 user/supervisor mode 00 IAC3 debug events can occur. 01 Reserved 10 IAC3 debug events can occur only if MSR[PR]=0. 11 IAC3 debug events can occur only if MSR[PR]=1.
50–51	IAC3ER	Instruction address compare 3 effective/real mode 00 IAC3 debug events are based on effective addresses. 01 IAC3 debug events are based on real addresses. 10 IAC3 debug events are based on effective addresses and can occur only if MSR[IS]=0. 11 IAC3 debug events are based on effective addresses and can occur only if MSR[IS]=1.

**Table 46. DBCR1 field descriptions (continued)**

Bits	Name	Description
52–53	IAC4US	Instruction address compare 4 user/supervisor mode 00 IAC4 debug events can occur. 01 Reserved 10 IAC4 debug events can occur only if MSR[PR]=0. 11 IAC4 debug events can occur only if MSR[PR]=1.
54–55	IAC4ER	Instruction address compare 4 effective/real mode 00 IAC4 debug events are based on effective addresses. 01 IAC4 debug events are based on real addresses. 10 IAC4 debug events are based on effective addresses and can occur only if MSR[IS]=0. 11 IAC4 debug events are based on effective addresses and can occur only if MSR[IS]=1.
56–57	IAC34M	Instruction address compare 3/4 mode 00 Exact address compare. IAC3 debug events can occur only if the instruction fetch address equals the value in IAC3. IAC4 debug events can occur only if the instruction fetch address equals the value in IAC4. 01 Address bit match. IAC3 and IAC4 debug events can occur only if the data storage access address, ANDed with the contents of IAC4, equals the value in IAC3, also ANDed with the contents of IAC4. If IAC3US≠IAC4US or IAC3ER≠IAC4ER, results are boundedly undefined. 10 Inclusive address range compare. IAC3 and IAC4 debug events can occur only if the instruction fetch address lies between the values specified in IAC3 and IAC4. If IAC3US≠IAC4US or IAC3ER≠IAC4ER, results are boundedly undefined. 11 Exclusive address range compare. IAC3 and IAC4 debug events can occur only if the instruction fetch address lies between the values specified in IAC3 and IAC4. If IAC3US≠IAC4US or IAC3ER≠IAC4ER, results are boundedly undefined.
58–63	—	Reserved, should be cleared.

**3.13.1.3 Debug control register 2 (DBCR2)**

DBCR2 is shown below.

**Figure 56. Debug control register 2 (DBCR2)**




Table 47. DBCR2 field descriptions

Bits	Name	Description
32–33	DAC1US	Data address compare 1 user/supervisor mode 00 DAC1 debug events can occur. 01 Reserved 10 DAC1 debug events can occur only if MSR[PR]=0. 11 DAC1 debug events can occur only if MSR[PR]=1.
34–35	DAC1ER	Data address compare 1 effective/real mode 00 DAC1 debug events are based on effective addresses. 01 DAC1 debug events are based on real addresses. 10 DAC1 debug events are based on effective addresses and can occur only if MSR[DS]=0. 11 DAC1 debug events are based on effective addresses and can occur only if MSR[DS]=1.
36–37	DAC2US	Data address compare 2 user/supervisor mode 00 DAC2 debug events can occur. 01 Reserved 10 DAC2 debug events can occur only if MSR[PR]=0. 11 DAC2 debug events can occur only if MSR[PR]=1.
38–39	DAC2ER	Data address compare 2 effective/real mode 00 DAC2 debug events are based on effective addresses. 01 DAC2 debug events are based on real addresses. 10 DAC2 debug events are based on effective addresses and can occur only if MSR[DS]=0. 11 DAC2 debug events are based on effective addresses and can occur only if MSR[DS]=1.
40–41	DAC12M	Data address compare 1/2 mode 00 Exact address compare. DAC1 debug events can occur only if the data access address equals the value in DAC1. DAC2 debug events can occur only if the data access address equals the value in DAC2. 01 Address bit match. DAC1 and DAC2 debug events can occur only if the data access address, ANDed with the contents of DAC2, equals the value in DAC1, also ANDed with the DAC2 contents. If DAC1US≠DAC2US or DAC1ER≠DAC2ER, results are boundedly undefined. 10 Inclusive address range compare. DAC1 and DAC2 debug events can occur only if the data access address lies between the values specified in DAC1 and DAC2. If DAC1US≠DAC2US or DAC1ER≠DAC2ER, results are boundedly undefined. 11 Exclusive address range compare. DAC1 and DAC2 debug events can occur only if the data access address lies between the values specified in DAC1 and DAC2. If DAC1US≠DAC2US or DAC1ER≠DAC2ER, results are boundedly undefined.
42	DAC1LNK	Data address compare 1 linked 0 No effect 1 DAC1 debug events are linked to IAC1 debug events. IAC1 debug events do not affect DBSR. When linked to IAC1, DAC1 debug events are conditioned based on whether the instruction also generated an IAC1 debug event.

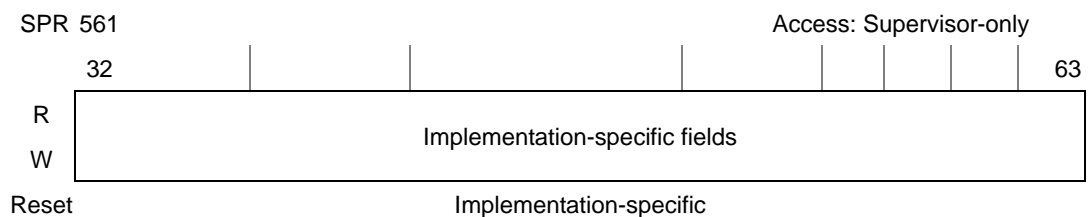
**Table 47. DBCR2 field descriptions (continued)**

Bits	Name	Description
43	DAC2LNK	Data address compare 2 linked 0 No effect 1 DAC 2 debug events are linked to IAC3 debug events. IAC3 debug events do not affect DBSR. When linked to IAC3, DAC2 debug events are conditioned based on whether the instruction also generated an IAC3 debug event. DAC2 can only be linked if DAC12M specifies exact address compare because DAC2 debug events are not generated in the other compare modes.
44–45	DVC1M	Data value compare 1 mode 00 DAC1 debug events can occur. 01 DAC1 debug events can occur only when all bytes in DBCR2[DVC1BE] in the data value of the data storage access match their corresponding bytes in DVC1. 10 DAC1 debug events can occur only when at least one of the bytes in DBCR2[DVC1BE] in the data value of the data storage access matches its corresponding byte in DVC1. 11 DAC1 debug events can occur only when all bytes in DBCR2[DVC1BE] within at least one of the half words of the data value of the data storage access match their corresponding bytes in DVC1.
46–47	DVC2M	Data value compare 2 mode 00 DAC2 debug events can occur. 01 DAC2 debug events can occur only when all bytes in DBCR2[DVC2BE] in the data value of the data storage access match their corresponding bytes in DVC2. 10 DAC2 debug events can occur only when at least one of the bytes in DBCR2[DVC2BE] in the data value of the data storage access matches its corresponding byte in DVC2. 11 DAC2 debug events can occur only when all bytes in DBCR2[DVC2BE] within at least one of the half words of the data value of the data storage access match their corresponding bytes in DVC2.
48–55	DVC1BE	Data value compare 1 byte enables. Specifies which bytes in the aligned data value being read or written by the storage access are compared to the corresponding bytes in DVC1.
56–63	DVC2BE	Data value compare 2 byte enables. Specifies which bytes in the aligned data value being read or written by the storage access are compared to the corresponding bytes in DVC2.

**3.13.1.4 Debug control register 3 (DBCR3)**

The debug APU defines the DBCR3, however its contents are implementation specific.


**Figure 57. Debug control register 3 (DBCR3)**



### 3.13.2 Debug status register (DBSR)

The DBSR, provides status debug events information for the most recent processor reset.

**Figure 58. Debug status register (DBSR)**



The DBSR is set through hardware, but is read through software using **mf spr** and cleared by writing ones to them; writing zeros has no effect.

**Table 48. DBSR field descriptions**

Bits	Name	Description												
32	IDE	Imprecise debug event. Set if MSR[DE] = 0 and a debug event causes its respective DBSR bit to be set.												
33	UDE	Unconditional debug event. Set if an unconditional debug event occurred. If the $\overline{UDE}$ signal (level sensitive, active low) is asserted, DBSR[UDE] is affected as follows: <table style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <tr> <td style="border: none;"><u>MSR[DE]</u></td> <td style="border: none;"><u>DBCRO[IDM]</u></td> <td style="border: none;">Action</td> </tr> <tr> <td style="border: none;">X</td> <td style="border: none;">0</td> <td style="border: none;">No action.</td> </tr> <tr> <td style="border: none;">0</td> <td style="border: none;">1</td> <td style="border: none;">DBSR[UDE] is set.</td> </tr> <tr> <td style="border: none;">1</td> <td style="border: none;">1</td> <td style="border: none;">DBSR[UDE] is set and a debug interrupt is taken.</td> </tr> </table>	<u>MSR[DE]</u>	<u>DBCRO[IDM]</u>	Action	X	0	No action.	0	1	DBSR[UDE] is set.	1	1	DBSR[UDE] is set and a debug interrupt is taken.
<u>MSR[DE]</u>	<u>DBCRO[IDM]</u>	Action												
X	0	No action.												
0	1	DBSR[UDE] is set.												
1	1	DBSR[UDE] is set and a debug interrupt is taken.												
34–35	MRR	Most recent reset. Set when a reset occurs. Undefined at power-up. See the implementation documentation.												
36	ICMP	Instruction complete debug event. Set if an instruction completion debug event occurred and DBCRO[ICMP] = 1.												
37	BRT	Branch taken debug event. Set if a branch taken debug event occurred (DBCRO[BRT]=1).												
38	IRPT	Interrupt taken debug event. Set if an interrupt taken debug event occurred (DBCRO[IRPT]=1).												
39	TRAP	Trap instruction debug event. Set if a trap Instruction debug event occurred (DBCRO[TRAP]=1).												
40	IAC1	Instruction address compare 1 debug event. Set if an IAC1 debug event occurred (DBCRO[IAC1]=1).												
41	IAC2	Instruction address compare 2 debug event. Set if an IAC2 debug event occurred (DBCRO[IAC2]=1).												


**Table 48. DBSR field descriptions (continued)**

Bits	Name	Description
42	IAC3	Instruction address compare 3 debug event. Set if an IAC3 debug event occurred (DBCR0[IAC3]=1).
43	IAC4	Instruction address compare 4 debug event. Set if an IAC4 debug event occurred (DBCR0[IAC4]=1).
44	DAC1R	Data address compare 1 read debug event. Set if a read-type DAC1 debug event occurred (DBCR0[DAC1]=10 or 11).
45	DAC1W	Data address compare 1 write debug event. Set if a write-type DAC1 debug event occurred (DBCR0[DAC1]=01 or 11).
46	DAC2R	Data address compare 2 read debug event. Set if a read-type DAC2 debug event occurred (DBCR0[DAC2]=10 or 11).
47	DAC2W	Data address compare 2 write debug event. Set if a write-type DAC2 debug event occurred (DBCR0[DAC2]=01 or 11).
48	RET	Return debug event. Set if a return debug event occurred (DBCR0[RET]=1).
49–56	—	Reserved, should be cleared.
57	CIRPT	Debug APU. Critical interrupt taken debug event. A critical interrupt taken debug event occurs when DBCR0[CIRPT] = 1 and a critical interrupt (any interrupt that uses the critical class, that is, uses CSRR0 and CSRR1) occurs. 0 No critical interrupt taken debug event has occurred. 1 A critical interrupt taken debug event occurred.
58	CRET	Debug APU. Critical interrupt return debug event. A critical interrupt return debug event occurs when DBCR0[CRET] = 1 and a return from critical interrupt (an <b>rfci</b> instruction is executed) occurs. 0 No critical interrupt return debug event has occurred. 1 A critical interrupt return debug event occurred.
59–63	—	Reserved, should be cleared.

### 3.13.3 Instruction address compare registers (IAC1–IAC4)

The instruction address compare registers (IAC1–IAC4) are each 64 bits, with bits 62–63 being reserved.

**Figure 59. Instruction address compare registers (IAC1–IAC4)**



A debug event may be enabled to occur upon an attempt to execute an instruction from an address specified in an IAC, inside or outside a range specified by IAC1 and IAC2 or, inside or outside a range specified by IAC3 and IAC4, or to blocks of addresses specified by the

combination of the IAC1 and IAC2, or to blocks of addresses specified by the combination of the IAC3 and IAC4. Because all instruction addresses are required to be word-aligned, the two low-order bits of the IACs are reserved and do not participate in the comparison to the instruction address.

### 3.13.4 Data address compare registers (DAC1–DAC2)

The data address compare registers (DAC1 and DAC2), are each 32 bits. A debug event may be enabled to occur upon loads, stores, or cache operations to an address specified in either DAC1 or DAC2, inside or outside a range specified by the DAC1 and DAC2, or to blocks of addresses specified by the combination of the DAC1 and DAC2.

**Figure 60. Data address compare registers (DAC1–DAC2)**




The contents of DAC1 or DAC2 are compared to the address generated by a data storage access instruction.

### 3.13.5 Data value compare registers (DVC1 and DVC2)

The data value compare registers (DVC1 and DVC2) are shown below. A DAC1R, DAC1W, DAC2R, or DAC2W debug event may be enabled to occur upon loads or stores of a specific data value specified in either or both of DVC1 and DVC2. DBCR2[DVC1M] and DBCR2[DVC1BE] control how the contents of DVC1 is compared with the value and DBCR2[DVC2M] and DBCR2[DVC2BE] control how the contents of DVC2 is compared with the value. [Table 47](#) describes the modes provided.

**Figure 61. Data value compare registers (DVC1–DVC2)**




## 3.14 SPE and SPFP APU registers

The SPE and SPFP include the signal processing and embedded floating-point status and control register (SPEFSCR), which is described in [Section 3.14.1: Signal processing, embedded floating-point status, control register \(SPEFSCR\)](#), and the SPE implements a 64-bit accumulator, described in [Section 3.14.2: Accumulator \(ACC\)](#).

### 3.14.1 Signal processing, embedded floating-point status, control register (SPEFSCR)

SPEFSCR, is used by the SPE and by the embedded floating-point APUs. 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 floating-point instructions affect only the low-element floating-point status flags and leave the high-element floating-point status flags undefined.

**Figure 62. Signal processing, embedded floating-point status and control register (SPEFSCR)**



**Table 49. SPEFSCR field descriptions**

Bits	Name	Description
32	SOVH	(SPE APU) Summary integer overflow high. Set when an SPE instruction sets OVH. This is a sticky bit that remains set until it is cleared by an <b>mtspr</b> instruction.
33	OVH	(SPE APU) Integer overflow high. Set when an overflow or underflow occurs in the upper word of the result of an SPE instruction.
34	FGH	(FP APUs) Embedded floating-point guard bit high. Used by the floating-point round interrupt handler. FGH is an extension of the low-order bits of the fractional result produced from a floating-point operation on the high word. FGH is zeroed if an overflow, underflow, or invalid input error is detected on the high element of a vector floating-point instruction. Execution of a scalar floating-point instruction leaves FGH undefined.
35	FXH	(SPFP APU) Embedded floating-point inexact bit high. Used by the floating-point round interrupt handler. FXH is an extension of the low-order bits of the fractional result produced from a floating-point operation on the high word. FXH represents the logical OR of all of the bits shifted right from the guard bit when the fractional result is normalized. FXH is zeroed if an overflow, underflow, or invalid input error is detected on the high element of a vector floating-point instruction. Execution of a scalar floating-point instruction leaves FXH undefined.



Table 49. SPEFSCR field descriptions (continued)

Bits	Name	Description
36	FINVH	(FP APUs) Embedded floating-point invalid operation/input error high. Set under any of the following conditions: Any operand of a high word vector floating-point instruction is Infinity, NaN, or Denorm The operation is a divide and the dividend and divisor are both 0 A conversion to integer or fractional value overflows. Execution of a scalar floating-point instruction leaves FINVH undefined.
37	FDBZH	(FP APUs) Embedded floating-point divide by zero high. Set when a vector floating-point divide instruction is executed with a divisor of 0 in the high word operand and the dividend is a finite non-zero number. Execution of a scalar floating-point instruction leaves FDBZH undefined.
38	FUNFH	(FP APUs) Embedded floating-point underflow high. Set when the execution of a vector floating-point instruction results in an underflow on the high word operation. Execution of a scalar floating-point instruction leaves FUNFH undefined.
39	FOVFH	(FP APUs) Embedded floating-point overflow high. Set when the execution of a vector floating-point instruction results in an overflow on the high word operation. Execution of a scalar floating-point instruction leaves FOVFH undefined.
40–41	—	Reserved, should be cleared.
42	FINXS	(FP APUs) Embedded floating-point inexact sticky flag. Set under the following conditions: – Execution of any scalar or vector floating-point instruction delivers an inexact result for either the low or high element and no floating-point data interrupt is taken for either element – 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 remains set until it is cleared by software.
43	FINVS	(FP APUs) Embedded floating-point invalid operation sticky flag. The sticky result of any floating-point instruction that causes FINVH or FINV to be set. That is, $FINVS \leftarrow FINVS \mid FINV \mid FINVH$ . This action may optionally be performed by hardware. To ensure proper operation, software should set this bit on the detection of FINV or FINVH set to one. FINVS remains set until it is cleared by software. <sup>(1)</sup>
44	FDBZS	(FP APUs) Embedded floating-point divide by zero sticky flag. Set when a floating-point divide instruction sets FDBZH or FDBZ. That is, $FDBZS \leftarrow FDBZS \mid FDBZH \mid FDBZ$ . FDBZS remains set until it is cleared by software.
45	FUNFS	(FP APUs) Embedded floating-point underflow sticky flag. Defined to be the sticky result of any floating-point instruction that causes FUNFH or FUNF to be set. That is, $FUNFS \leftarrow FUNFS \mid FUNF \mid FUNFH$ . This action may optionally be performed by hardware. To ensure proper operation, software should set this bit on the detection of FUNF or FUNFH being set. FUNFS remains set until it is cleared by software. <sup>1</sup>
46	FOVFS	(FP APUs) Embedded floating-point overflow sticky flag. defined to be the sticky result of any floating-point instruction that causes FOVH or FOVF to be set. That is, $FOVFS \leftarrow FOVFS \mid FOVF \mid FOVFH$ . This action may optionally be performed by hardware. To ensure proper operation, software should set this bit on the detection of FOVF or FOVFH being set. FOVFS remains set until it is cleared by software. <sup>1</sup>

Table 49. SPEFSCR field descriptions (continued)

Bits	Name	Description
47	MODE	(FP APUs) Embedded floating-point operating mode. Controls the operating mode of the embedded floating-point operations defined in the SPE, and the embedded floating-point APUs. 0 Default hardware results operating mode 1 Reserved.
48	SOV	(SPE APU) Summary integer overflow low. Set when an SPE instruction sets OV. This sticky bit remains set until an <b>mtspr</b> writes a 0 to this bit.
49	OV	(SPE APU) Integer overflow low. OV is set when an overflow or underflow occurs in the lower word of the result of an SPE instruction.
50	FG	(FP APUs) Embedded floating-point guard bit (low/scalar) Used by the floating-point round interrupt handler. FG is an extension of the low-order bits of the fractional result produced from a floating-point operation on the low word or any scalar floating-point operation. FG is cleared if an overflow, underflow, or invalid input error is detected on either the low element of a vector floating-point instruction or any scalar floating-point instruction.
51	FX	(FP APUs) Embedded floating-point inexact bit (low/scalar). Used by the floating-point round interrupt handler. FX is an extension of the low-order bits of the fractional result produced from a floating-point operation on the low word or any scalar floating-point instruction. FX represents the logical OR of all of 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 either the low element of a vector floating-point instruction or any scalar floating-point instruction.
52	FINV	(FP APUs) Embedded floating-point invalid operation/input error (low/scalar). Set by the following conditions: – Any operand of a low-word vector or scalar floating-point operation is Infinity, NaN, or Denorm – The operation is a divide and the dividend and divisor are both 0 – A conversion to integer or fractional value overflows
53	FDBZ	(FP APUs) Embedded floating-point divide by zero (low/scalar). Set when a scalar or vector floating-point divide instruction is executed with a divisor of 0 in the low word operand and the dividend is a finite non-zero number.
54	FUNF	(FP APUs) Embedded floating-point underflow (low/scalar). Set when execution of a scalar or vector floating-point instruction results in an underflow on the low word operation.
55	FOVF	(FP APUs) Embedded floating-point overflow (low/scalar). Set when the execution of a scalar or vector floating-point instruction results in an overflow on the low word operation.
56	—	Reserved, should be cleared.
57	FINXE	(FP APUs) Embedded floating-point round (inexact) exception enable 0 Exception disabled 1 Exception enabled. A floating-point round interrupt is taken if no other interrupt is taken, and if FG   FGH   FX   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	(FP APUs) Embedded floating-point invalid operation/input error exception enable 0 Exception disabled 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FINV or FINVH.

Table 49. SPEFSCR field descriptions (continued)

Bits	Name	Description
59	FDBZE	(FP APUs) Embedded floating-point divide by zero exception enable 0 Exception disabled 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FDBZ or FDBZH.
60	FUNFE	(FP APUs) Embedded floating-point underflow exception enable 0 Exception disabled 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FUNF or FUNFH.
61	FOVFE	(FP APUs) Embedded floating-point overflow exception enable 0 Exception disabled 1 Exception enabled. A floating-point data interrupt is taken if a floating-point instruction sets FOVF or FOVFH.
62–63	FRMC	(FP APUs) Embedded floating-point rounding mode control 00 Round to Nearest 01 Round toward Zero 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. 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 hardware that manages this sticky bit by performing an operation on a NaN and observing whether hardware sets this sticky bit. In the absence of doing this, if it desired that software written will work on all processors that support embedded floating-point, software should check the appropriate status bits and set the sticky bit itself (if hardware also performs this operation, the action is redundant).

### 3.14.2 Accumulator (ACC)

The 64-bit architectural accumulator register holds the results of the multiply accumulate (MAC) forms of SPE integer instructions. The accumulator 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. The accumulator, however, has to be explicitly cleared when starting a new MAC loop. Based upon the type of instruction, an accumulator can hold either a single 64-bit value or a vector of two 32-bit elements.


The Initialize Accumulator instruction (**evmra**) is provided to initialize the accumulator. This instruction is described in [Chapter 7: Instruction set](#).

## 3.15 Alternate time base registers (ATBL and ATBU)

The alternate time base counter (ATB), is formed by concatenating the upper and lower alternate time base registers (ATBU and ATBL). ATBL (SPR 526) provides read-only access to the 64-bit alternate time base counter, which is incremented at an implementation-defined frequency. ATB registers are accessible in both user and supervisor mode.

Like the TB implementation, ATBL is an aliased name for ATB.

**Figure 63. Alternate time base register lower (ATBL)**




**Table 50. ATBL field descriptions**

Bits	Name	Description
32–63	ATBCL	Alternate time base counter lower. Lower 32 bits of the alternate time base counter

The ATBU register, provides read-only access to the upper 32 bits of the alternate time base counter. It is accessible in both user and supervisor mode.

**Figure 64. Alternate time base register upper (ATBU)**



**Table 51. ATBU field descriptions**

Bits	Name	Description
32–63	ATBCU	Alternate time base counter upper. Upper 32 bits of the alternate time base counter

### 3.16 Performance monitor registers (PMRs)

The EIS defines a set of register resources used exclusively by the performance monitor. PMRs are similar to the SPRs defined in the Book E architecture and are accessed by **mtpmr** and **mfpmr**, which are also defined by the EIS. [Table 52](#) lists supervisor-level PMRs. User-level software that attempts to read or write supervisor-level PMRs causes a privilege exception.

**Table 52. Performance monitor registers—supervisor level**

Abbreviation	Register name	PMR number	pmr[0–4]	pmr[5–9]	Section/page
PMGC0	Performance monitor global control register 0	400	01100	10000	<a href="#">Section 3.16.1 on page 126</a>

**Table 52. Performance monitor registers—supervisor level (continued)**

Abbreviation	Register name	PMR number	pmr[0–4]	pmr[5–9]	Section/page
PMLCa0	Performance monitor local control a0	144	00100	10000	<a href="#">Section 3.16.3 on page 127</a>
PMLCa1	Performance monitor local control a1	145	00100	10001	
PMLCa2	Performance monitor local control a2	146	00100	10010	
PMLCa3	Performance monitor local control a3	147	00100	10011	
PMLCb0	Performance monitor local control b0	272	01000	10000	<a href="#">Section 3.16.5 on page 128</a>
PMLCb1	Performance monitor local control b1	273	01000	10001	
PMLCb2	Performance monitor local control b2	274	01000	10010	
PMLCb3	Performance monitor local control b3	275	01000	10011	
PMC0	Performance monitor counter 0	16	00000	10000	<a href="#">Section 3.16.7 on page 130</a>
PMC1	Performance monitor counter 1	17	00000	10001	
PMC2	Performance monitor counter 2	18	00000	10010	
PMC3	Performance monitor counter 3	19	00000	10011	

User-level PMRs in [Table 53](#) are read-only and are accessed with `mfpmr`. Attempting to write user-level registers in supervisor or user mode causes an illegal instruction exception.

**Table 53. Performance monitor registers—user level (read-only)**


Abbreviation	Register name	PMR number	pmr[0–4]	pmr[5–9]	Section/page
UPMGC0	User performance monitor global control register 0	384	01100	00000	<a href="#">Section 3.16.3 on page 127</a>
UPMLCa0	User performance monitor local control a0	128	00100	00000	<a href="#">Section 3.16.4 on page 128</a>
UPMLCa1	User performance monitor local control a1	129	00100	00001	
UPMLCa2	User performance monitor local control a2	130	00100	00010	
UPMLCa3	User performance monitor local control a3	131	00100	00011	
UPMLCb0	User performance monitor local control b0	256	01000	00000	<a href="#">Section 3.16.6 on page 130</a>
UPMLCb1	User performance monitor local control b1	257	01000	00001	
UPMLCb2	User performance monitor local control b2	258	01000	00010	
UPMLCb3	User performance monitor local control b3	259	01000	00011	
UPMC0	User performance monitor counter 0	0	00000	00000	<a href="#">Section 3.16.7 on page 130</a>
UPMC1	User performance monitor counter 1	1	00000	00001	
UPMC2	User performance monitor counter 2	2	00000	00010	
UPMC3	User performance monitor counter 3	3	00000	00011	

### 3.16.1 Global control register 0 (PMGC0)

The performance monitor global control register (PMGC0), controls all performance monitor counters.

#### 3.16.1.1 Performance monitor global control register 0 (PMGC0)/ User performance monitor global control register 0 (UPMGC0)

**Figure 65. Global control register 0 (PMGC0)**



PMGC0 is cleared by a hard reset. Reading this register does not change its contents.

**Table 54. PMGC0 field descriptions**

Bits	Name	Description
32	FAC	Freeze all counters. When FAC is set by hardware or software, PMLCx[FC] maintains its current value until it is changed by software. 0 The PMCs are incremented (if permitted by other PM control bits). 1 The PMCs are not incremented.
33	PMIE	Performance monitor interrupt enable 0 Performance monitor interrupts are disabled. 1 Performance monitor interrupts are enabled and occur when an enabled condition or event occurs.
34	FCECE	Freeze counters on enabled condition or event 0 The PMCs can be incremented (if permitted by other PM control bits). 1 The PMCs can be incremented (if permitted by other PM control bits) only until an enabled condition or event occurs. When an enabled condition or event occurs, PMGC0[FAC] is set. It is up to software to clear FAC.
35–50	—	Reserved, should be cleared.
51–52	TBSEL	Time base selector. Selects the time base bit that can cause a time base transition event (the event occurs when the selected bit changes from 0 to 1). 00 TB[63] (TBL[31]) 01 TB[55] (TBL[23]) 10 TB[51] (TBL[19]) 11 TB[47] (TBL[15])  Time base transition events can be used to periodically collect information about processor activity. In multiprocessor systems in which TB registers are synchronized among processors, time base transition events can be used to correlate the performance monitor data obtained by the several processors. For this use, software must specify the same TBSEL value for all processors in the system. Because the time-base frequency is implementation-dependent, software should invoke a system service program to obtain the frequency before choosing a value for TBSEL.

**Table 54. PMGC0 field descriptions (continued)**

Bits	Name	Description
53–54	—	Reserved, should be cleared.
55	TBEE	Time base transition event exception enable 0 Exceptions from time base transition events are disabled. 1 Exceptions from time base transition events are enabled. A time base transition is signalled to the performance monitor if the TB bit specified in PMGC0[TBSEL] changes from 0 to 1. Time base transition events can be used to freeze the counters (PMGC0[FCECE]) or signal an exception (PMGC0[PMIE]). Changing PMGC0[TBSEL] while PMGC0[TBEE] is enabled may cause a false 0 to 1 transition that signals the specified action (freeze, exception) to occur immediately. Although the interrupt signal condition may occur with MSR[EE] = 0, the interrupt cannot be taken until MSR[EE] = 1.
55–63	—	Reserved, should be cleared.

**3.16.2 User global control register 0 (UPMGC0)**

The contents of PMGC0 are reflected to UPMGC0, which is read by user-level software. UPMGC0 is read with the `mfpmr` instruction using PMR384.


**3.16.3 Local control A registers (PMLCa0–PMLCa3)**

The local control A registers 0–3 (PMLCa0–PMLCa3), function as event selectors and give local control for the corresponding performance monitor counters. PMLCa works with the corresponding PMLCb register.

**3.16.3.1 Local control A registers (PMLCa0–PMLCa3)/  
User local control A registers (UPMLCa0–UPMLCa3)**

**Figure 66. Local control A registers (PMLCa0–PMLCa3)**

PMLCa0 (PMR144)	UPMLCa0 (PMR128)	Access: PMLCa0–PMLCa3: supervisor-only
PMLCa1 (PMR145)	UPMLCa1 (PMR129)	UPMLCa0–UPMLCa3: supervisor/user read-only
PMLCa2 (PMR146)	UPMLCa2 (PMR130)	
PMLCa3 (PMR147)	UPMLCa3 (PMR131)	



Reset

All zeros

Table 55. PMLCa0–PMLCa3 field descriptions

Bits	Name	Description
32	FC	Freeze counter 0 The PMC is incremented (if permitted by other PM control bits). 1 The PMC is not incremented.
33	FCS	Freeze counter in supervisor state 0 The PMC is incremented (if permitted by other PM control bits). 1 The PMC is not incremented if MSR[PR] = 0.
34	FCU	Freeze counter in user state 0 The PMC is incremented (if permitted by other PM control bits). 1 The PMC is not incremented if MSR[PR] = 1.
35	FCM1	Freeze counter while mark = 1 0 The PMC is incremented (if permitted by other PM control bits). 1 The PMC is not incremented if MSR[PMM] = 1.
36	FCM0	Freeze counter while mark = 0 0 The PMC is incremented (if permitted by other PM control bits). 1 The PMC is not incremented if MSR[PMM] = 0.
37	CE	Condition enable 0 PMCx overflow conditions cannot occur. (PMCx cannot cause interrupts, cannot freeze counters.) 1 Overflow conditions occur when the most-significant-bit of PMCx is equal to one.  It is recommended that CE be cleared when counter PMCx is selected for chaining.
38–40	—	Reserved, should be cleared.
41–47	EVENT	Event selector. Up to 128 events selectable.
48–63	—	Reserved, should be cleared.

### 3.16.4 User local control A registers (UPMLCa0–UPMLCa3)

The contents of PMLCa0–PMLCa3 are reflected to UPMLCa0–UPMLCa3, which are read by user-level software with `mfpmr` using PMR numbers in [Table 53](#).


### 3.16.5 Local control B registers (PMLCb0–PMLCb3)

Local control B registers (PMLCb0–PMLCb3), specify a threshold value and a multiple to apply to a threshold event selected for the corresponding performance monitor counter. PMLCb works with the corresponding PMLCa.



### 3.16.5.1 Local control B registers (PMLCb0–PMLCb3)/User local control B registers (UPMLCb0–UPMLCb3)

**Figure 67. Local control B registers (PMLCb0–PMLCb3)**



**Table 56. PMLCb0 –PMLCb3 field descriptions**

Bits	Name	Description
32–52	—	Reserved, should be cleared.
53–55	THRESHMUL	Threshold multiple 000 Threshold field is multiplied by 1 (PMLCb $\eta$ [THRESHOLD] * 1) 001 Threshold field is multiplied by 2 (PMLCb $\eta$ [THRESHOLD] * 2) 010 Threshold field is multiplied by 4 (PMLCb $\eta$ [THRESHOLD] * 4) 011 Threshold field is multiplied by 8 (PMLCb $\eta$ [THRESHOLD] * 8) 100 Threshold field is multiplied by 16 (PMLCb $\eta$ [THRESHOLD] * 16) 101 Threshold field is multiplied by 32 (PMLCb $\eta$ [THRESHOLD] * 32) 110 Threshold field is multiplied by 64 (PMLCb $\eta$ [THRESHOLD] * 64) 111 Threshold field is multiplied by 128 (PMLCb $\eta$ [THRESHOLD] * 128)
56–57	—	Reserved, should be cleared.
58–63	THRESHOLD	Threshold. Only events that exceed this value are counted. Events to which a threshold value applies are implementation-dependent as are the dimension (for example duration in cycles) and the granularity with which the threshold value is interpreted. By varying the threshold value, software can profile event characteristics. For example, if PMC1 is configured to count cache misses that last longer than the threshold value, software can obtain the distribution of cache miss durations for a given program by monitoring the program repeatedly using a different threshold value each time.

### 3.16.6 User local control B registers (UPMLCb0–UPMLCb3)


The contents of PMLCb0–PMLCb3 are reflected to UPMLCb0–UPMLCb3, which are read by user-level software with **mfpmr** using the PMR numbers in [Table 53](#).

### 3.16.7 Performance monitor counter registers (PMC0–PMC3)

The performance monitor counter registers PMC0–PMC3, are 32-bit counters that can be programmed to generate interrupt signals when they overflow. Each counter is enabled to count 128 events.

#### 3.16.7.1 Performance monitor counter registers (PMC0–PMC3)/User performance monitor counter registers (UPMC0–UPMC3)

**Figure 68. Performance monitor counter registers (PMC0–PMC3)**



**Table 57. PMC0–PMC3 field descriptions**

Bits	Name	Description
32	OV	Overflow. When this bit is set, it indicates this counter reaches its maximum value.
33–63	Counter Value	Indicates the number of occurrences of the specified event.

Counters overflow when the high-order bit (the sign bit) becomes set; that is, they reach the value 2,147,483,648 (0x8000\_0000). However, an exception is not signaled unless PMGC0[PMIE] and PMLCax[CE] are also set as appropriate.

The interrupts are masked by clearing MSR[EE]. An interrupt that is signaled while MSR[EE] is zero is not taken until MSR[EE] is set. Setting PMGC0[FCECE] forces counters to stop counting when an enabled condition or event occurs.

Software is expected to use **mtpmr** to explicitly set PMCs to non-overflowed values. Setting an overflowed value may cause an erroneous exception. For example, if both PMGC0[PMIE] and PMLCax[CE] are set and the **mtpmr** loads an overflowed value into PMCx, an interrupt may be generated without an event counting having taken place.

PMC registers are accessed with **mtpmr** and **mfpmr** using the PMR numbers in [Table 52](#).

### 3.16.8 User performance monitor counter registers (UPMC0–UPMC3)

The contents of PMC0–PMC3 are reflected to UPMC0–UPMC3, which are read by user-level software with the **mfpmr** instruction using the PMR numbers in [Table 53](#).

### 3.17 Device control registers (DCRs)

Book E defines the existence of a DCR address space and the instructions to access them, but does not define particular DCRs. The on-chip DCRs exist architecturally outside the processor core and thus are not part of Book E.

DCRs may control the use of on-chip peripherals, such as memory controllers (specific DCR definitions would be provided in the implementation's user's manual).

The contents of DCR DCRN can be read into a GPR using `mfdcr rD,DCRN`. GPR contents can be written into DCR DCRN using `mtdcr DCRN,rS`.

If DCRs are implemented, they are described as part of the implementation documentation.

### 3.18 Book E SPR model

This section describes SPR invalid references, synchronization requirements, and preserved, reserved, and allocated registers.

#### 3.18.1 Invalid SPR references

System behavior when an invalid SPR is referenced depends on the privilege level.

- If the invalid SPR is accessible in user mode ( $\text{SPR}[5] = 0$ ), an illegal instruction exception is taken.
- If the invalid SPR is accessible only in supervisor mode ( $\text{SPR}[5] = 1$ ) and the core complex is in supervisor mode ( $\text{MSR}[\text{PR}] = 0$ ), the results of the attempted access are boundedly undefined.
- If the invalid SPR address is accessible only in supervisor mode (bit 5 of an SPR number = 1) and the core complex is not in supervisor mode ( $\text{MSR}[\text{PR}] = 1$ ), a privilege exception is taken. These results are summarized in [Table 58](#).

**Table 58. System response to an invalid spr reference**

SPR address bit 5	MSR[PR]	Response
0 (User)	x	Illegal exception
1 (Supervisor)	0 (Supervisor)	Boundedly undefined
	1 (User)	Privilege exception

#### 3.18.2 Synchronization requirements for SPRs

Synchronization requirements for accessing certain SPRs are shown in [Table 59](#). Except for these SPRs, there are no synchronization requirements for accessing SPRs beyond those stated in Book E. (Note that requirements may be different for different implementations.)

Table 59. Synchronization requirements for sprs

Registers	Instruction	Instruction required before	Instruction required after
DBCR0	mtspr docr0	None	isync
DBCR1	mtspr docr1	None	isync
HID0	mtspr hid0	None	isync
HID1	mtspr hid1	None	isync
L1CSR0	mtspr l1csr0	<b>msync, isync</b>	isync
L1CSR1	mtspr l1csr1	None	isync
MAS $n$	mtspr mas $n$	None	isync
MMUCSR0	mtspr mmucsr0	None	isync
PID $n$	mtspr pid $n$	None	isync
SPEFSCR	mtspr spefscr	None	isync

### 3.18.3 Reserved SPRs

An undefined SPR number in the range 0x000–0x1FF (0–511) that is not preserved is reserved.

### 3.18.4 Allocated SPRs

SPR numbers allocated for implementation-dependent use are 0x200–0x3FF (512–1023).

Table 60. Allocated SPRs defined by the EIS

SPR	Mnemonic	Register
48	PID0 <sup>(1)</sup>	Process ID register 0. This is not truly an allocated SPR; however, Book E defines only this PID register and refers to it as PID rather than PID0.
512	SPEFSCR	Signal processing and embedded floating-point status and control register
515	L1CFG0	L1 cache configuration register 0
516	L1CFG1	L1 cache configuration register 1
528	IVOR32	SPE APU unavailable exception
529	IVOR33	Embedded floating-point data exception
530	IVOR34	Embedded floating-point round exception
531	IVOR35	Performance monitor Interrupt vector offset register
570	MCSRR0	Machine-check save/restore register 0
571	MCSRR1	Machine-check save/restore register 1
572	MCSR	Machine check syndrome register
573	MCAR	Machine check address register
624	MAS0	MMU assist register 0
625	MAS1	MMU assist register 1

Table 60. Allocated SPRs defined by the EIS (continued)

SPR	Mnemonic	Register
626	MAS2	MMU assist register 2
627	MAS3	MMU assist register 3
628	MAS4	MMU assist register 4
629	MAS5	MMU assist register 5
630	MAS6	MMU assist register 6
633	PID1	Process ID register 1
634	PID2	Process ID register 2
...	PID $n$	Additional PID registers may be defined in this space
688	TLB0CFG	TLB configuration register 0
689	TLB1CFG	TLB configuration register 1
944	MAS7	MMU assist register 7
1008	HID0	Hardware implementation dependent register 0
1009	HID1	Hardware implementation dependent register 1
1010	L1CSR0	L1 cache control and status register 0
1011	L1CSR1	L1 cache control and status register
1012	MMUCSR0	MMU control and status register 0
1015	MMUCFG	MMU configuration register
1023	SVR	System version register

1. An update to a PID register must always be followed by an **isync**.

## 4 Instruction model

The architecture specifications allow for different processor implementations, which may provide extensions to or deviations from the architectural descriptions. This chapter provides information about the Book E architecture and the Book E implementation standards (EIS), which defines auxiliary processing units (APUs) and other architectural extensions that define additional instructions, registers, and interrupts. For more information, see [Chapter 8: Auxiliary processing units \(APUs\)](#).

### 4.1 Operand conventions

This section describes operand conventions as they are represented in the Book E architecture. These conventions follow the basic descriptions in the classic PowerPC architecture with some changes in terminology. For example, distinctions between user and supervisor-level instructions are maintained, but the designations—UISA, VEA, and OEA—do not apply. Detailed descriptions are provided of conventions used for storing values in registers and memory, accessing processor registers, and representing data in these registers.

#### 4.1.1 Data organization in memory and data transfers

Bytes in memory are numbered consecutively starting with 0. Each number is the address of the corresponding byte.

Memory operands can be bytes, half words, words, or double words or, for the load/store multiple instruction type and load/store string instructions, a sequence of bytes or words. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Operand length is implicit for each instruction.

#### 4.1.2 Alignment and misaligned accesses

The operand of a single-register memory access instruction has an alignment boundary equal to its length. An operand's address is misaligned if it is not a multiple of its width.

The concept of alignment is also applied more generally to data in memory. For example, a 12-byte data item is said to be word-aligned if its address is a multiple of four.

Some instructions require their memory operands to have certain alignment. In addition, alignment can affect performance. For single-register memory access instructions, the best performance is obtained when memory operands are aligned.

Instructions are 32 bits (one word) long and must be word-aligned. Note, however, that the VLE extension provides both 16- and 32-bit instructions.

See [Section 4.6.6.10: VLE instruction alignment and byte ordering](#).

[Table 61](#) lists characteristics for memory operands for single-register memory access instructions.

Table 61. Address characteristics of aligned operands

Operand	Operand Length	Addr[60–63] if Aligned
Byte (or string)	8 bits	xxxx <sup>(1)</sup>
Half word	2 bytes	xxx0
Word	4 bytes	xx00
Double word	8 bytes	x000

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

Note that **lmw**, **stmw**, **lwarx**, and **stwcx**. instructions that are not word aligned cause an alignment exception.

## 4.2 Instruction set summary

Instructions are divided into the following functional categories:

- Integer instructions—These include arithmetic and logical instructions. See [Section 4.3.1.1: Integer instructions](#).
- Floating-point instructions—These include floating-point vector and scalar arithmetic instructions. See [Section 4.6.1.2: Embedded vector and scalar floating-point APU instructions](#). Note that some implementations do not support Book E–defined floating-point instructions or registers.
- Load and store instructions—See [Section 4.3.1.8: Load and store instructions](#).
- Flow control instructions—These include branching instructions, CR logical instructions, trap instructions, and other instructions that affect the instruction flow. See [Section 4.3.1.11: Branch and flow control instructions](#).
- Processor control instructions—These instructions are used for synchronizing memory accesses. See [Section 4.6.6.1: Processor control instructions](#).
- Memory synchronization instructions—These instructions are used for memory synchronizing. See [Section 4.3.1.15: Memory synchronization instructions](#).
- Memory control instructions—These instructions provide control of caches and TLBs. See [Section 4.3.1.17: Memory control instructions](#), and [Section 4.3.2.2: Supervisor-level memory control instructions](#).
- Signal processing instructions—These include a set of vector arithmetic and logic instructions optimized for signal processing. See [Section 4.6.1: SPE and embedded floating-point APUs](#).

*Note:* *Instruction groupings used here do not indicate the execution unit that processes a particular instruction or group of instructions. This information, which is useful for scheduling instructions most effectively, is provided in the execution chapter for the implementation.*

Integer instructions operate on word operands. Book E floating-point instructions operate on single-precision and double-precision floating-point operands. The PowerPC architecture uses instructions that are 4 bytes long and word-aligned. It provides for byte, half-word, and word operand loads and stores between memory and a set of 32 general-purpose registers (GPRs). It provides for word and double-word operand loads and stores between memory and a set of 32 floating-point registers (FPRs).

Arithmetic and logical instructions do not read or modify memory. To use the contents of a memory location in a computation and then modify the same or another location, the memory contents must be loaded into a register, modified, and then written to the target location using load and store instructions.

The description of each instruction includes the mnemonic and a formatted list of operands. To simplify assembly language programming, a set of simplified mnemonics and symbols is provided for some of the frequently used instructions; see [Appendix B: Simplified mnemonics for PowerPC instructions](#), for a complete list of simplified mnemonics. Programs written to be portable across the various assemblers for the PowerPC architecture should not assume the existence of mnemonics not described in that document.

## 4.2.1 Classes of instructions

Instructions belong to one of the following four classes:

- Defined instructions (See [Section 4.2.1.2: Defined instruction class](#))
- Allocated instructions (See [Section 4.2.1.3: Allocated instruction class](#))
- Preserved instructions (See [Section 4.2.1.4: Preserved instruction class](#))
- Reserved (illegal or no-op) instructions (See [Section 4.2.1.5: Reserved instruction class](#))

The class is determined by examining the primary opcode and any extended opcode. If the opcode, or combination of opcode and extended opcode, is not that of a defined, allocated, preserved, or reserved instruction, the instruction is illegal.

### 4.2.1.1 Definition of boundedly undefined

If instructions are encoded with incorrectly set bits in reserved fields, the results on execution can be said to be boundedly undefined. If a user-level program executes the incorrectly coded instruction, the resulting undefined results are bounded in that a spurious change from user to supervisor state is not allowed, and the level of privilege exercised by the program in relation to memory access and other system resources cannot be exceeded. Boundedly undefined results for a given instruction can vary between implementations and between execution attempts in the same implementation.

### 4.2.1.2 Defined instruction class

This class of instructions consists of all the instructions defined in Book E. In general, defined instructions are guaranteed to be supported within a Book E system as specified by the architecture, either within the processor implementation itself or within emulation software supported by the system operating software.



For implementations that only provide the 32-bit subset of Book E, emulation of the 64-bit behavior of the defined instructions is not supported. See [Appendix D: Guidelines for 32-bit book E](#).

Any attempt to execute a defined instruction results in one of the following events:

- An illegal instruction exception-type program interrupt, if an implementation does not recognize the instruction
- An unimplemented instruction exception-type program interrupt, if the instruction is recognized but not supported by the implementation and is not a floating-point instruction
- An unimplemented instruction exception-type program interrupt, if the instruction is recognized but not supported by the implementation, and is a floating-point instruction and floating-point processing is enabled
- The floating-point unavailable interrupt if the instruction is recognized but is not supported by the implementation or is a floating-point instruction and floating-point processing is disabled
- The floating-point unavailable interrupt when floating-point processing is disabled and a floating-point instruction is recognized and is not supported by the implementation
- If an instruction is recognized and supported by the implementation, the processor performs the actions described in the rest of this document. The architected behavior may cause other exceptions.

A defined instruction may be retained by future versions of Book E as a defined instruction, or may be reclassified as a preserved instruction (process of removal from the architecture) and eventually classified as reserved-illegal.

#### 4.2.1.3 Allocated instruction class

This class of instructions contains the set of instructions (a set of primary opcodes, as well as a set of extended opcodes for certain primary opcodes) used for implementation-specific instructions. [Table 62](#) lists blocks of opcodes allocated for implementation-dependent use.

**Table 62. Allocated instructions**

Primary opcode	Extended opcodes
0	All instruction encodings (bits 6–31) except 0x0000_0000 <sup>(1)</sup> .
4	All instruction encodings (bits 6–31) SPE and embedded floating-point instructions
19	Extended opcodes (bits 21–30) 0buuuuu_0u11u
31	Extended opcodes (bits 21–30) uuuuu_0u11u
59	Extended opcodes (bits 21–30) uuuuu_0u10u
63	Extended opcodes (bits 21–30) uuuuu_0u10u (except 00000_01100 <b>frsp</b> )

1. Instruction encoding 0x0000\_0000 is and always will be reserved-illegal.

Allocated instructions are allocated to purposes that are outside the scope of Book E for implementation-dependent and application-specific use.

Any attempt to execute an allocated instruction results in one of the following:

- An illegal instruction exception-type program interrupt, if the instruction is not recognized by the implementation
- An unimplemented instruction exception-type program interrupt, if the instruction is recognized and enabled for execution but the implementation does not support direct execution of the instruction. This option may be used to allow emulation for unsupported allocated instructions.
- A floating-point unavailable interrupt, if an allocated instruction that extends the floating-point capabilities is recognized and floating-point processing is disabled
- If an allocated instruction is implemented, the processor performs the actions described in the user’s manual. Implementation-dependent behavior may cause other exceptions.

An allocated instruction is guaranteed by Book E to remain allocated.

*Note: Some allocated instructions may have associated new process state, and, therefore, may provide an associated enable bit, similar in function to MSR[FP] for floating-point instructions. For such instructions, being enabled for execution implies that any associated enable bit is set to allow, or enable, instruction execution. For such instructions, the architecture provides an auxiliary processor unavailable interrupt vector in case execution of such an instruction is attempted when execution is disabled. For example, MSR[SPE] enables the SPE unavailable interrupt. Other allocated instructions may not have any associated new state and therefore may not require an associated enable bit. If supported by an implementation, such instructions are assumed to be always enabled for execution.*

**4.2.1.4 Preserved instruction class**

The preserved instruction class supports backward compatibility with the PowerPC architecture. An attempt to execute a preserved instruction results in one of the following:

- If the implementation does not recognize the instruction, an illegal instruction exception-type program interrupt occurs.
- If the instruction is recognized and supported by the implementation, the processor performs the actions defined in the previous version of the architecture.

Future versions of Book E may retain a preserved instruction as a preserved instruction, may reclassify it as an allocated instruction, or may adopt it as part of Book E.

Preserved opcodes are listed in [Table 63](#).

**Table 63. Preserved instructions**

Primary opcode	Extended opcodes
0	No preserved extended opcodes
4	No preserved extended opcodes
19	No preserved extended opcodes



Table 63. Preserved instructions (continued)

Primary opcode	Extended opcodes
31	Extended opcodes (bits 21–30) 210 0b00110_10010 ( <b>mtsr</b> ) 242 0b00111_10010 ( <b>mtsrin</b> ) 370 0b01011_10010 ( <b>tlbia</b> ) 306 0b01001_10010 ( <b>tlbie</b> ) 371 0b01011_10011 ( <b>mftb</b> ) 595 0b10010_10011 ( <b>mfsr</b> ) 659 0b10100_10011 ( <b>mfsrin</b> ) 310 0b01001_10110 ( <b>eciwx</b> ) 438 0b01101_10110 ( <b>ecowx</b> )
59	No preserved extended opcodes
63	No preserved extended opcodes

#### 4.2.1.5 Reserved instruction class

This class of instructions consists of all instruction primary opcodes (and associated extended opcodes, if applicable) that do not belong to either the defined, allocated, or preserved instruction classes.

Reserved instructions are available for future extensions of Book E. That is, some future version of Book E may define any of these instructions to perform new functions or make them available for implementation-dependent use as allocated instructions. There are two types of reserved instructions, reserved-illegal and reserved-nop.

Attempts to execute a reserved-illegal instruction cause an illegal instruction exception-type program interrupt (see [Section 5.7.6: Alignment interrupt](#)) on implementations conforming to the current version of Book E. Reserved-illegal instructions are, therefore, available for future extensions to Book E that would affect architected state. Such extensions might include new forms of integer or floating-point arithmetic or new forms of load or store instructions that write their result in an architected register.

Attempts to execute a reserved-nop instruction either do not affect implementations conforming to the current version of Book E (that is, treated as a no-operation instruction), or cause an illegal instruction exception-type program interrupt (see [Section 5.7.7: Program interrupt](#)). Reserved-nop instructions are available for future architecture extensions that do not affect architected state. Such extensions might include performance-enhancing hints such as new forms of cache touch instructions and could be added while remaining functionally compatible with implementations of previous versions of Book E

A reserved-illegal instruction may be retained by future versions of Book E as a reserved-illegal instruction, may be subsequently reclassified as an allocated instruction, or may even be employed in the role of a subsequently defined instruction.

A reserved-nop instruction may be retained by future versions of Book E as a reserved-nop instruction, may be subsequently reclassified as an allocated instruction, or may even be employed in the role of a subsequently defined instruction that has no effect on architected state.

## 4.2.2 Instruction forms

This section describes preferred instruction forms, addressing modes, and synchronization.

### 4.2.2.1 Preferred instruction forms (no-op)

The Or Immediate (**ori**) instruction has the following preferred form for expressing a no-op:

```
ori 0,0,0
```

### 4.2.2.2 Invalid instruction forms

Some of the defined instructions have invalid forms. An instruction form is invalid if one or more fields of the instruction, excluding the opcode field(s), are coded incorrectly in a manner that can be deduced by examining only the instruction encoding.

Attempts to execute an invalid form of an instruction either causes an illegal instruction type program interrupt or yields boundedly undefined results. Any exceptions to this rule are stated in the instruction descriptions.

Some kinds of invalid form instructions can be deduced just from examining the instruction layout. These are listed below.

- Field shown as reserved but coded as nonzero
- Field shown as containing a particular value but coded as some other value

These invalid forms are not discussed further.

Other invalid instruction forms can be deduced by detecting an invalid encoding of one or more of the instruction operand fields. These kinds of invalid form are identified in the instruction descriptions.

- Branch conditional and branch conditional extended instructions (undefined encoding of BO field)
- Load with update instructions ( $rD = rA$  or  $rA = 0$ )
- Store with update instructions ( $rA = 0$ )
- Load multiple instruction ( $rA$  or  $rB$  in range of registers to be loaded)
- Load string immediate instructions ( $rA$  in range of registers to be loaded)
- Load string indexed instructions ( $rD = rA$  or  $rD = rB$ )
- Load/store floating-point with update instructions ( $rA = 0$ )

## 4.2.3 Addressing modes

This section describes conventions for addressing memory and for calculating effective addresses (EAs) as defined by the Book E architecture for 32-bit implementations.

### 4.2.3.1 Memory addressing

A program references memory using the effective address computed by the processor when it executes a memory access or branch instruction (or other instructions as described in [Section 4.3.1.18: User-level cache instructions](#), and [Section 4.3.2.3: Supervisor-level cache instruction](#), or when it fetches the next sequential instruction.

### 4.2.3.2 Memory operands

Bytes in memory are numbered consecutively starting with 0. Each number is the address of the corresponding byte.

Memory operands may be bytes, half words, words or, for the load/store multiple and load/store string instructions, a sequence of words or bytes. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Byte ordering can be either big endian or little endian (see [Section 4.2.3.4: Byte ordering](#)). The default byte and bit ordering is big endian.

Operand length is implicit for each instruction with respect to memory alignment. The operand of a scalar memory access instruction has a natural alignment boundary equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is said to be misaligned. For more information about alignment, see [Section 4.1.2: Alignment and misaligned accesses](#).

### 4.2.3.3 Effective address calculation

The 32-bit address computed by the processor when executing a memory access or branch instruction (or certain other instructions described in [Section 4.3.1.18: User-level cache instructions](#), [Section 4.3.2.3: Supervisor-level cache instruction](#), and [Section 4.3.2.4: Supervisor-level tlb management instructions](#)), or when fetching the next sequential instruction, is called the effective address (EA) and specifies a byte in memory. For a memory access instruction, if the sum of the EA and the operand length exceeds the maximum EA, the memory access is considered to be undefined.

Effective address arithmetic, except for next sequential instruction address computations, wraps around from the maximum address,  $2^{32}-1$ , to address 0.

Data memory addressing modes

Book E supports the following data memory addressing modes:

- Base+displacement addressing mode—The 16-bit D field is sign-extended and added to the contents of the GPR designated by rA or to zero if rA = 0. Instructions that use this addressing mode are of the D instruction format.
- Base+index addressing mode—The contents of the GPR designated by rB (or the value 0 for **lswi** and **stswi**) are added to the contents of the GPR designated by rA or to zero if rA = 0. Instructions that use this addressing mode are of the X instruction format.
- Base+displacement extended addressing mode—The 12-bit DE field is sign-extended and added to the contents of the GPR designated by rA or to zero if rA = 0 to produce the 32-bit EA. Instructions that use this addressing mode are of the DE instruction format.
- Base+displacement extended scaled addressing mode—The 12-bit DES field is concatenated on the right with zeros, sign-extended, and added to the contents of the GPR designated by rA or to zero if rA = 0 to produce the 32-bit EA. Instructions that use this addressing mode are of the DES instruction format.

In addition, APUs may provide additional addressing modes.

#### 4.2.3.3.1 Instruction memory addressing modes

Instruction memory addressing modes correspond with instructions forms, as follows:

- I-form branch instructions—The 24-bit LI field is concatenated on the right with 0b00, sign-extended, and then added to either the address of the branch instruction if AA = 0, or to 0 if AA = 1.
- Taken B-form branch instructions—The 14-bit BD field is concatenated on the right with 0b00, sign-extended, and then added to either the address of the branch instruction if AA = 0, or to 0 if AA = 1.
- Taken XL-form branch instructions—The contents of bits LR[32–61] or CR[32–61] are concatenated on the right with 0b00.
- Sequential instruction fetching (or non-taken branch instructions)—The value 4 is added to the address of the current instruction to form the 32-bit EA of the next instruction. If the address of the current instruction is 0xFFFF\_FFFC, the address of the next sequential instruction is undefined.
- Any branch instruction with LK = 1—The value 4 is added to the address of the current instruction and the 32-bit result is placed into the LR. If the address of the current instruction is 0xFFFF\_FFFC, the result placed into the LR is undefined.

Although some implementations may support next sequential instruction address computations wrapping from the highest address 0xFFFF\_FFFC to 0x0000\_0000 as part of the instruction flow, users are strongly encouraged not to depend on this behavior. Doing so can reduce the portability of their software. If code must span this boundary, software should place a non-linking branch at address 0xFFFF\_FFFC, which always branches to address 0x0000\_0000 (either absolute or relative branches work).

See also [Appendix D: Guidelines for 32-bit book E](#).

#### 4.2.3.4 Byte ordering

If scalars (individual data items and instructions) were indivisible, there would be no such concept as byte ordering. It is meaningless to consider the order of bits or groups of bits within the smallest addressable unit of memory, because nothing can be observed about such order. Only when scalars, which the programmer and processor regard as indivisible quantities, can comprise more than one addressable unit of memory does the question of order arise.

For a machine in which the smallest addressable unit of memory is the 64-bit double word, there is no question of the ordering of bytes within double words. All transfers of individual scalars between registers and memory are of double words, and the address of the byte containing the high-order 8 bits of a scalar is no different from the address of a byte containing any other part of the scalar.

For Book E, as for most computer architectures currently implemented, the smallest addressable unit of memory is the 8-bit byte. Many scalars are half words and words (double words in 64-bit implementations) which consist of groups of bytes. When a word-length scalar is moved from a register to memory, the scalar occupies four consecutive byte addresses. It thus becomes meaningful to discuss the order of the byte addresses with respect to the value of the scalar: which byte contains the highest-order eight bits of the scalar, which byte contains the next-highest-order 8 bits, and so on.

Given a scalar that contains multiple bytes, the choice of byte ordering is essentially arbitrary. There are  $4! = 24$  ways to specify the ordering of 4 bytes within a word but only two of these orderings are sensible:

- The ordering that assigns the lowest address to the highest-order (left-most) 8 bits of the scalar, the next sequential address to the next-highest-order eight bits, and so on. This ordering is called big endian because the big (most-significant) end of the scalar, considered as a binary number, comes first in memory. The 68000 is an example of a processor using this byte ordering.
- The ordering that assigns the lowest address to the lowest-order (right-most) 8 bits of the scalar, the next sequential address to the next-lowest-order eight bits, and so on. This ordering is called little endian because the little (least-significant) end of the scalar, considered as a binary number, comes first in memory. The Intel 8086 is an example of a processor using this byte ordering.

Book E provides support for both big- and little-endian byte ordering in the form of a memory attribute. See [Section 6.4.8: Permission attributes](#), and [Section 6.2.1: Memory/Cache access attributes](#).

#### 4.2.3.5 Synchronization requirements

This section describes synchronization requirements for special registers and TLBs. Changing the value in certain system registers and invalidating TLB entries can have the side effect of altering the context in which data addresses and instruction addresses are interpreted, and in which instructions are executed. For example, changing  $MSR[IS] = 0$  to and  $MSR[IS] = 1$  has the side effect of changing address space. Such effects need not occur in program order (program order refers to the execution of instructions in the strict order in which they occur in the program), and therefore may require explicit synchronization by software.

An instruction that alters the context in which data addresses or instruction addresses are interpreted, or in which instructions are executed, is called context altering. This section covers all such context-altering instructions. The required software synchronization for each is shown in [Table 64](#).

The notation 'CSI' in the tables means any context-synchronizing instruction (such as, **sc**, **isync**, **rftci**, or **rfti**). A context-synchronizing interrupt (that is, any interrupt except non-recoverable machine check) can be used instead of a context-synchronizing instruction. If it is, phrases like 'the synchronizing instruction,' below, should be interpreted as meaning the instruction at which the interrupt occurs. If no software synchronization is required before (after) a context-altering instruction, "the synchronizing instruction before (after) the context-altering instruction" should be interpreted as meaning the context-altering instruction itself.

The synchronizing instruction before the context-altering instruction ensures that all instructions up to and including that synchronizing instruction are fetched and executed in the context that existed before the alteration. The synchronizing instruction after the context-altering instruction ensures that all instructions after that synchronizing instruction are fetched and executed in the context established by the alteration. Instructions after the first synchronizing instruction, up to and including the second synchronizing instruction, may be fetched or executed in either context.

If a sequence of instructions contains context-altering instructions and contains no instructions that are affected by any of the context alterations, no software synchronization is required within the sequence.

Note: Sometimes advantage can be taken of the fact that certain instructions that occur naturally in the program, such as the **rfi/rfci** at the end of an interrupt handler, provide the required synchronization.

No software synchronization is required before altering the MSR (except perhaps when altering the WE bit: see the tables), because **mtmsr** is execution synchronizing. No software synchronization is required before most of the other alterations shown in the [Section : Instruction fetch and/or execution](#) section in [Table 64](#), because all instructions before the context-altering instruction are fetched and decoded before the context-altering instruction executes (the processor must determine whether any of the preceding instructions are context synchronizing)

[Table 64](#) identifies the software synchronization requirements for data access for all context-altering instructions.

**Table 64. Synchronization requirements**

Context altering instruction or event	Required before	Required after	Notes
<b>Data Accesses</b>			
interrupt	None	None	
<b>mtmsr</b> (DS)	None	CSI	
<b>mtmsr</b> (ME)	None	CSI	(1)
<b>mtmsr</b> (PR)	None	CSI	
<b>mtspr</b> (DAC1, DAC2, DVC1, DVC2)	—	—	(2)
<b>mtspr</b> (DBCR0, DBCR2)	—	—	(2)
<b>mtspr</b> (DBSR)	—	—	(2)
<b>mtspr</b> (PID <sub>n</sub> )	CSI	CSI	
<b>rfci</b>	None	None	
<b>rfi</b>	None	None	
<b>sc</b>	None	None	
<b>tlbivax</b>	CSI	CSI or <b>msync</b>	(3),(4)
<b>tlbwe</b>	CSI	CSI or <b>msync</b>	(3),(4)
<b>Instruction fetch and/or execution</b>			
Interrupt	None	None	
<b>mtmsr</b> (CE)	None	None	(5)
<b>mtmsr</b> (DE)	None	CSI	
<b>mtmsr</b> (EE)	None	None	(3)
<b>mtmsr</b> (FE0)	None	CSI	
<b>mtmsr</b> (FE1)	None	CSI	
<b>mtmsr</b> (FP)	None	CSI	
<b>mtmsr</b> (IS)	None	CSI	(6)
<b>mtmsr</b> (ME)	None	CSI	(1)
<b>mtmsr</b> (PR)	None	CSI	



**Table 64. Synchronization requirements (continued)**

Context altering instruction or event	Required before	Required after	Notes
<b>mtmsr</b> (WE)	—	—	(7)
<b>mtspr</b> (DBCR0, DBCR1)	—	—	(2)
<b>mtspr</b> (DBSR)	—	—	(2)
<b>mtspr</b> (DEC)	None	None	(8)
<b>mtspr</b> (IAC1, IAC2, IAC3, IAC4)	—	—	(2)
<b>mtspr</b> (IVORi)	None	None	
<b>mtspr</b> (IVPR)	None	None	
<b>mtspr</b> (PID)	None	CSI	(6)
<b>mtspr</b> (TCR)	None	None	(8)
<b>mtspr</b> (TSR)	None	None	(8)
<b>rftci</b>	None	None	
<b>rfti</b>	None	None	
<b>sc</b>	None	None	
<b>tlbivax</b>	None	CSI or <b>msync</b>	(3),(4)
<b>tlbwe</b>	None	CSI or <b>msync</b>	(3),(4)
<b>wrtee</b>	None	None	(5)
<b>wrteei</b>	None	None	(5)

1. A context synchronizing instruction is required after altering MSR[ME] to ensure that the alteration takes effect for subsequent machine check interrupts, which may not be recoverable and therefore may not be context synchronizing.
2. Synchronization requirements for changing any of the debug registers are implementation-dependent and are specified in the user's manual for the implementation.
3. For data accesses, the context synchronizing instruction before the **tlbwe** or **tlbivax** instruction ensures that all storage accesses due to preceding instructions have completed to a point at which they have reported all exceptions they cause.  
 The context synchronizing instruction after the **tlbwe** or **tlbivax** ensures that subsequent storage accesses (data and instruction) use the updated value in any affected TLB entries. It does not ensure that all storage accesses previously translated by the TLB entries being updated have completed with respect to storage; if these completions must be ensured, the **tlbwe** or **tlbivax** must be followed by an **msync** instruction as well as by a context synchronizing instruction.  
 The following sequence shows why it is necessary for data accesses to ensure that all storage accesses due to instructions before a **tlbwe** or **tlbivax** have completed to a point at which they have reported all exceptions they will cause. Assume that valid TLB entries exist for the target storage location when the sequence starts.
  1. A program issues a load or store to a page.
  2. The same program executes a **tlbwe** or **tlbivax** that invalidates the corresponding TLB entry.
  3. The load or store instruction finally executes, and gets a TLB miss exception. The TLB miss exception is semantically incorrect. In order to prevent it, a context synchronizing instruction must be executed between steps 1 and 2.
4. Multiprocessor systems have other requirements to synchronize what is called 'TLB shoot down' (that is, to invalidate one or more TLB entries on all processors in the multiprocessor system and to be able to determine that the invalidations have completed and that all side effects of the invalidations have taken effect).

5. The effect of changing MSR[EE] or MSR[CE] is immediate.  
 If an **mtmsr**, **wrtee**, or **wrteei** clears MSR[EE], an external input, decremter, or fixed-interval timer interrupt does not occur after the instruction executes.  
 If an **mtmsr**, **wrtee**, or **wrteei** changes MSR[EE] from 0 to 1 when an external input, decremter, fixed-interval timer, or higher priority enabled exception exists, the corresponding interrupt occurs immediately after the **mtmsr**, **wrtee**, or **wrteei** executes and before the next instruction is executed in the program that sets MSR[EE].  
 If an **mtmsr** clears MSR[CE], a critical input, or watchdog timer interrupt does not occur after the instruction is executed.  
 If an **mtmsr** changes MSR[CE] from 0 to 1 when a critical input, watchdog timer, or higher priority enabled exception exists, the corresponding interrupt occurs immediately after **mtmsr** executes, and before the next instruction is executed in the program that set MSR[CE].
6. The alteration must not cause an implicit branch in real address space. Thus the real address of the context-altering instruction and of each subsequent instruction, up to and including the next context synchronizing instruction, must be independent of whether the alteration has taken effect.
7. Synchronization requirements for changing the wait state enable are implementation-dependent, and are specified in the user's manual for the implementation.
8. The elapsed time between the decremter reaching zero, or the transition of the selected time base bit for the fixed-interval timer or the watchdog timer, and the signalling of the decremter, fixed-interval timer or the watchdog timer exception is not defined.

#### 4.2.3.6 Context synchronization

An instruction or event is context synchronizing if it satisfies the requirements listed below. Context-synchronizing operations include instructions **isync**, **sc**, **rfi**, **rfdi**, and **rfmci**, and most interrupts.

1. The operation is not initiated or, in the case of **isync**, does not complete until all instructions already in execution have completed to a point at which they have reported all exceptions they cause.
2. The instructions that precede the operation complete execution in the context (including such parameters as privilege level, address space, and memory protection) in which they were initiated.
3. If the operation directly causes an interrupt (for example, **sc** directly causes a system call interrupt) or is an interrupt, the operation is not initiated until no interrupt-causing exception exists having higher priority than the exception associated with the interrupt. See [Section 5.11: Exception priorities](#).
4. The instructions that follow the operation are fetched and executed in the context established by the operation as required by the sequential execution model. (This requirement dictates that any prefetched instructions be discarded and that any effects and side effects of executing them speculatively may also be discarded, except as described in [Section 6.2.2.1: Memory access ordering](#).)

A context-synchronizing operation is necessarily execution synchronizing. Unlike **msync** and **mbar**, such operations do not affect the order of memory accesses with respect to other mechanisms.

#### 4.2.3.7 Execution synchronization

An instruction is execution synchronizing if it satisfies items <sup>(1)</sup> and <sup>(2)</sup> of the definition of context synchronization. **msync** is treated like **isync** with respect to item <sup>(1)</sup> (that is, the conditions described in item <sup>(1)</sup> apply to completion of **msync**). Execution synchronizing instructions include **msync**, **mtmsr**, **wrtee**, and **wrteei**. All context-synchronizing instructions are execution synchronizing.

Unlike a context-synchronizing operation, an execution synchronizing instruction need not ensure that the instructions following it execute in the context established by that execution synchronizing instruction. This new context becomes effective sometime after the execution

synchronizing instruction completes and before or at a subsequent context-synchronizing operation.

#### Instruction-related interrupts

Interrupts are caused either directly by the execution of an instruction or by an asynchronous event. In either case, an exception may cause one of several types of interrupts to be invoked.

Examples of interrupts that can be caused directly by the execution of an instruction include but are not limited to the following:

- An attempt to execute a reserved-illegal instruction (illegal instruction exception-type program interrupt)
- An attempt by an application program to execute a privileged instruction (privileged instruction exception-type program interrupt)
- An attempt by an application program to access a privileged SPR (privileged instruction exception-type program interrupt)
- An attempt by an application program to access an SPR that does not exist (unimplemented operation instruction exception-type program interrupt)
- An attempt by a system program to access an SPR that does not exist (boundedly undefined)
- Execution of a defined instruction using an invalid form (illegal instruction exception-type program interrupt, unimplemented operation exception-type program interrupt, or privileged instruction exception-type program interrupt)
- An attempt to access a memory location that is either unavailable (instruction TLB error interrupt or data TLB error interrupt) or not permitted (instruction storage interrupt or data storage interrupt)
- An attempt to access memory with an EA alignment not supported by the implementation (alignment interrupt)
- Execution of a **system call** instruction (system call interrupt)
- Execution of a **trap** instruction whose trap condition is met (trap type program interrupt)
- Execution of a floating-point instruction when floating-point instructions are unavailable (floating-point unavailable interrupt)
- Execution of a floating-point instruction that causes a floating-point enabled exception to exist (floating-point enabled exception-type program interrupt)
- Execution of a defined instruction that is not implemented by the implementation (illegal instruction exception or unimplemented operation exception-type program interrupt)
- Execution of an allocated instruction that is not implemented by the implementation (illegal instruction exception or unimplemented operation exception-type program interrupt)
- Execution of an allocated instruction when the auxiliary instruction is unavailable (auxiliary processor unavailable interrupt).
- Execution of an allocated instruction that causes an auxiliary enabled exception (enabled exception-type program interrupt).

APUs, such as the SPE, may define additional instruction-caused exceptions and interrupts. The invocation of an interrupt is precise, except that if one of the imprecise modes for invoking the floating-point enabled exception-type program interrupt is in effect the invocation of the floating-point enabled exception-type program interrupt may be imprecise. When the interrupt is invoked imprecisely, the excepting instruction does not appear to

complete before the next instruction starts (because one of the effects of the excepting instruction, namely the invocation of the interrupt, has not yet occurred).

[Chapter 5: Interrupts and exceptions](#) describes interrupt conditions in detail.

### 4.3 Instruction set overview

This section provides a brief overview of the Book E and Book E instructions.

*Note:* Some instructions have the following optional features:

- *CR update*—The dot (.) suffix on the mnemonic enables the update of the CR.
- *Overflow option*—The o suffix indicates that the overflow bit in the XER is enabled.

#### 4.3.1 Book E user-level instructions

This section discusses the user-level instructions defined in the Book E architecture.

##### 4.3.1.1 Integer instructions

This section describes the integer instructions. These consist of the following:

- Integer arithmetic instructions
- Integer compare instructions
- Integer logical instructions
- Integer rotate and shift instructions

Integer instructions use the content of the GPRs as source operands and place results into GPRs and the XER and CR fields.

##### 4.3.1.2 Integer arithmetic instructions

[Table 65](#) lists the integer arithmetic instructions for the PowerPC processors.

**Table 65. Integer arithmetic instructions**

Name	Mnemonic	Syntax
Add	<b>add</b> ( <b>add. addo addo.</b> )	rD,rA,rB
Add carrying	<b>addc</b> ( <b>addc. addco addco.</b> )	rD,rA,rB
Add extended	<b>adde</b> ( <b>adde. addeo addeo.</b> )	rD,rA,rB
Add immediate	<b>addi</b>	rD,rA,SIMM
Add immediate carrying	<b>addic</b>	rD,rA,SIMM
Add immediate carrying and record	<b>addic.</b>	rD,rA,SIMM
Add immediate shifted	<b>addis</b>	rD,rA,SIMM
Add to minus one extended	<b>addme</b> ( <b>addme. addmeo addmeo.</b> )	rD,rA
Add to zero extended	<b>addze</b> ( <b>addze. addzeo addzeo.</b> )	rD,rA
Divide word	<b>divw</b> ( <b>divw. divwo divwo.</b> )	rD,rA,rB
Divide word unsigned	<b>divwu</b> ( <b>divwu. divwuo divwuo.</b> )	rD,rA,rB
Multiply high word	<b>mulhw</b> ( <b>mulhw.</b> )	rD,rA,rB

**Table 65. Integer arithmetic instructions (continued)**

Name	Mnemonic	Syntax
Multiply high word unsigned	<b>mulhwu</b> (mulhwu.)	rD,rA,rB
Multiply low immediate	<b>mulli</b>	rD,rA,SIMM
Multiply low word	<b>mullw</b> (mullw. mullwo mullwo.)	rD,rA,rB
Negate	<b>neg</b> (neg. nego nego.)	rD,rA
Subtract from	<b>subf</b> (subf. subfo subfo.)	rD,rA,rB
Subtract from carrying	<b>subfc</b> (subfc. subfco subfco.)	rD,rA,rB
Subtract from extended	<b>subfe</b> (subfe. subfeo subfeo.)	rD,rA,rB
Subtract from immediate carrying	<b>subfic</b>	rD,rA,SIMM
Subtract from minus one extended	<b>subfme</b> (subfme. subfmeo subfmeo.)	rD,rA
Subtract from zero extended	<b>subfze</b> (subfze. subfzeo subfzeo.)	rD,rA

Although there is no subtract immediate instruction, its effect can be achieved by using an **addi** instruction with the immediate operand negated. Simplified mnemonics are provided that include this negation. Subtract instructions subtract the second operand (rA) from the third operand (rB). Simplified mnemonics are provided in which the third operand is subtracted from the second. See [Appendix B: Simplified mnemonics for PowerPC instructions](#), for examples.

According to Book E, an implementation that executes instructions with the overflow exception enable bit (OE) set or that sets the carry bit (CA) can either execute these instructions slowly or prevent execution of the subsequent instruction until the operation completes. The summary overflow (SO) and overflow (OV) bits in the XER are set to reflect an overflow condition of a 32-bit result only if the instruction's OE bit is set.

#### 4.3.1.2.1 Integer compare instructions

The integer compare instructions algebraically or logically compare the contents of register rA with either the **zero-extended value of the UIMM** operand, the **sign-extended value of the SIMM** operand, or the contents of rB. The comparison is signed for **cmpi** and **cmp** and unsigned for **cmpli** and **cmpl**. [Table 66](#) lists integer compare instructions. Note that the L bit must be 0 for 32-bit implementations.

**Table 66. Integer 32-Bit compare instructions (L = 0)**

Name	Mnemonic	Syntax
Compare	<b>cmp</b>	crD,L,rA,rB
Compare immediate	<b>cmpi</b>	crD,L,rA,SIMM
Compare logical	<b>cmpl</b>	crD,L,rA,rB
Compare logical immediate	<b>cmpli</b>	crD,L,rA,UIMM

The crD operand can be omitted if the result of the comparison is to be placed in CR0. Otherwise the target CR field must be specified in crD by using an explicit field number.

For information on simplified mnemonics for the integer compare instructions see [Appendix B: Simplified mnemonics for PowerPC instructions.](#)

### 4.3.1.2.2 Integer logical instructions

The logical instructions shown in [Table 67](#) perform bit-parallel operations on the specified operands. Logical instructions with the CR updating enabled (uses dot suffix) and instructions **andi.** and **andis.** set CR field CR0 to characterize the result of the logical operation. Logical instructions do not affect XER[SO], XER[OV], or XER[CA].

See [Appendix B](#), for simplified mnemonic examples for integer logical operations.

**Table 67. Integer logical instructions**

Name	Mnemonic	Syntax	Implementation notes
AND	and (and.)	rA,rS,rB	—
AND Immediate	andi.	rA,rS,UIM M	—
AND Immediate Shifted	andis.	rA,rS,UIM M	—
AND with Complement	<b>andc</b> (andc.)	rA,rS,rB	—
Count Leading Zeros Word	<b>cntlzw</b> (cntlzw.)	rA,rS	—
Equivalent	<b>eqv</b> (eqv.)	rA,rS,rB	—
Extend Sign Byte	<b>extsb</b> (extsb.)	rA,rS	—
Extend Sign Half Word	<b>extsh</b> (extsh.)	rA,rS	—
NAND	<b>nand</b> (nand.)	rA,rS,rB	—
NOR	<b>nor</b> (nor.)	rA,rS,rB	—
OR	<b>or</b> (or.)	rA,rS,rB	—
OR Immediate	ori	rA,rS,UIM M	Book E defines <b>ori r0,r0,0</b> as the preferred form for a no-op. The dispatcher may discard this instruction and dispatch it only to the completion queue but not to any execution unit.
OR Immediate Shifted	oris	rA,rS,UIM M	—
OR with Complement	orc ( <b>orc.</b> )	rA,rS,rB	—
XOR	<b>xor</b> (xor.)	rA,rS,rB	—
XOR Immediate	xori	rA,rS,UIM M	—
XOR Immediate Shifted	xoris	rA,rS,UIM M	—

### 4.3.1.2.3 Integer rotate and shift instructions

Rotation operations are performed on data from a GPR, and the result, or a portion of the result, is returned to a GPR. Integer rotate instructions, summarized in [Table 68](#), rotate the

contents of a register. The result is either inserted into the target register under control of a mask (if a mask bit is set the associated bit of the rotated data is placed into the target register, and if the mask bit is cleared the associated bit in the target register is unchanged) or ANDed with a mask before being placed into the target register. [Appendix B: Simplified mnemonics for PowerPC instructions](#), lists simplified mnemonics that allow simpler coding of often used functions such as clearing the left- or right-most bits of a register, left or right justifying an arbitrary field, and simple rotates and shifts.

**Table 68. Integer rotate instructions**

Name	Mnemonic	Syntax
Rotate left word Immediate then AND with mask	<b>rlwinm</b> (rlwinm.)	rA,rS,SH,MB,ME
Rotate left word then AND with mask	<b>rlwnm</b> (rlwnm.)	rA,rS,rB,MB,ME
Rotate left word Immediate then mask insert	<b>rlwimi</b> (rlwimi.)	rA,rS,SH,MB,ME

The integer shift instructions ([Table 69](#)) perform left and right shifts. Immediate-form logical (unsigned) shift operations are obtained by specifying masks and shift values for certain rotate instructions. Simplified mnemonics (shown in [Appendix B: Simplified mnemonics for PowerPC instructions](#)) are provided to simplify coding of such shifts.

Multiple-precision shifts can be programmed as shown in [C.2: Multiple-precision shifts](#). The integer shift instructions are summarized in [Table 69](#).

**Table 69. Integer shift instructions**

Name	Mnemonic	Syntax
Shift Left Word	slw (slw.)	rA,rS,rB
Shift Right Word	srw (srw.)	rA,rS,rB
Shift Right Algebraic Word Immediate	srawi (srawi.)	rA,rS,SH
Shift Right Algebraic Word	sraw (sraw.)	rA,rS,rB

### 4.3.1.3 Floating-point instructions

This section describes the floating-point instructions as they are defined by Book E.

The rules followed in assigning new primary and extended opcodes.

- Primary opcode 63 is used for the double-precision arithmetic instructions as well as miscellaneous instructions (for example, FPSCR manipulation instructions). Primary opcode 59 is used for the single-precision arithmetic instructions.
- The single-precision instructions for which there is a corresponding double-precision instruction have the same format and extended opcode as that double-precision instruction.
- In assigning new extended opcodes for primary opcode 63, the following regularities are maintained. In addition, all new X-form instructions in primary opcode 63 have bits 21–22 = 11.
  - Bit 26 = 1 if and only if the instruction is A-form.
  - Bits 26–29 = 0b0000 if and only if the instruction is a comparison or **mcrfs** (if and only if the instruction sets an explicitly designated CR field).

- Bits 26–28 = 0b001 if and only if the instruction explicitly refers to or sets the FPSCR (that is, is an FPSCR instruction) and is not **mcrfs**.
- Bits 26–30 = 0b01000 if and only if the instruction is a move register instruction, or any other instruction that does not refer to or set the FPSCR.
- In assigning extended opcodes for primary opcode 59, the following regularities have been maintained. They are based on those rules for primary opcode 63 that apply to the instructions having primary opcode 59. In particular, primary opcode 59 has no FPSCR instructions, so the corresponding rule does not apply.
  - If there is a corresponding instruction with primary opcode 63, its extended opcode is used.
  - Bit 26 = 1 if and only if the instruction is A form.
  - Bits 26–30 = 0b01000 if and only if the instruction is a move register instruction, or any other instruction that does not refer to or set the FPSCR.

#### 4.3.1.3.1 Floating-point load instructions

There are two basic forms of load instruction: single-precision and double-precision. Because the FPRs support only floating-point double format, single-precision load floating-point instructions convert single-precision data to double format prior to loading the operand into the target FPR. The conversion and loading steps are as follows.

Let  $WORD_{0:31}$  be the floating-point single-precision operand accessed from memory.

##### Normalized Operand

```

if  $WORD_{1:8} > 0$  and  $WORD_{1:8} < 255$  then
   $FPR(\mathbf{frD})_{0:1} \leftarrow WORD_{0:1}$ 
   $FPR(\mathbf{frD})_2 \leftarrow \neg WORD_1$ 
   $FPR(\mathbf{frD})_3 \leftarrow \neg WORD_1$ 
   $FPR(\mathbf{frD})_4 \leftarrow \neg WORD_1$ 
   $FPR(\mathbf{frD})_{5:63} \leftarrow WORD_{2:31} \parallel 2^9 0$ 

```

##### Denormalized Operand

```

if  $WORD_{1:8} = 0$  and  $WORD_{9:31} \neq 0$  then
  sign  $\leftarrow WORD_0$ 
  exp  $\leftarrow -126$ 
   $frac_{0:52} \leftarrow 0b0 \parallel WORD_{9:31} \parallel 2^9 0$ 
  normalize the operand
  do while  $frac_0 = 0$ 
     $frac \leftarrow frac_{1:52} \parallel 0b0$ 
    exp  $\leftarrow exp - 1$ 
   $FPR(\mathbf{frD})_0 \leftarrow sign$ 
   $FPR(\mathbf{frD})_{1:11} \leftarrow exp + 1023$ 
   $FPR(\mathbf{frD})_{12:63} \leftarrow frac_{1:52}$ 

```

##### Zero/Infinity/NaN

```

if  $WORD_{1:8} = 255$  or  $WORD_{1:31} = 0$  then
   $FPR(\mathbf{frD})_{0:1} \leftarrow WORD_{0:1}$ 
   $FPR(\mathbf{frD})_2 \leftarrow WORD_1$ 
   $FPR(\mathbf{frD})_3 \leftarrow WORD_1$ 
   $FPR(\mathbf{frD})_4 \leftarrow WORD_1$ 
   $FPR(\mathbf{frD})_{5:63} \leftarrow WORD_{2:31} \parallel 2^9 0$ 

```

For double-precision load floating-point instructions, conversion is not required because the data from memory is copied directly into the FPR.



Many floating-point load instructions have an update form, in which GPR(*rA*) is updated with the EA. For these forms, if  $rA \neq 0$  and  $rA \neq rD$ , the EA is placed into GPR(*rA*) and the memory element (byte, half word, word, or double word) addressed by EA is loaded into FPR(*rD*). If  $rA=0$  or  $rA=rD$ , the instruction form is invalid.

Floating-point load accesses cause a data storage interrupt if the program is not allowed to read the location. Floating-point load memory accesses cause a data TLB error interrupt if the program attempts to access memory that is unavailable. The floating-point load instruction set is shown in [Table 70](#).

**Table 70. Floating-point load instruction set**

Instruction	Mnemonic	Syntax
Load Floating-Point Double	<b>lfd</b>	frD,D( <i>rA</i> )
Load Floating-Point Double with Update	<b>lfd</b>	frD,D( <i>rA</i> )
Load Floating-Point Double Extended	<b>lfde</b>	frD,DES( <i>rA</i> )
Load Floating-Point Double with Update Extended	<b>lfdue</b>	frD,DES( <i>rA</i> )
Load Floating-Point Double Indexed	<b>lfdx</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Double with Update Indexed	<b>lfd</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Double Indexed Extended	<b>lfdxe</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Double with Update Indexed Extended	<b>lfd</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Single	<b>lfs</b>	frD,D( <i>rA</i> )
Load Floating-Point Single with Update	<b>lfs</b>	frD,D( <i>rA</i> )
Load Floating-Point Single Extended	<b>lfse</b>	frD,DES( <i>rA</i> )
Load Floating-Point Single with Update Extended	<b>lfsue</b>	frD,DES( <i>rA</i> )
Load Floating-Point Single Indexed	<b>lfsx</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Single with Update Indexed	<b>lfs</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Single Indexed Extended	<b>lfsxe</b>	frD, <i>rA</i> , <i>rB</i>
Load Floating-Point Single with Update Indexed Extended	<b>lfsue</b>	frD, <i>rA</i> , <i>rB</i>

#### 4.3.1.3.2 Floating-point store instructions

There are three basic forms of store instruction: single-precision, double-precision, and integer. The integer form is provided by the optional store floating-point as integer word instruction (**stfiwx**), described in [Chapter 7: Instruction set](#). Because the FPRs support only floating-point double format for floating-point data, single-precision store floating-point instructions convert double-precision data to single-precision format before storing the operand. The conversion steps are as follows.

Let  $WORD_{0:31}$  be the word in memory written to.

**No Denormalization Required (includes Zero / Infinity / NaN)**

if  $FPR(FRS)_{1:11} > 896$  or  $FPR(FRS)_{1:63} = 0$  then

$WORD_{0:1} \leftarrow FPR(FRS)_{0:1}$

$WORD_{2:31} \leftarrow FPR(FRS)_{5:34}$

```

Denormalization Required
if  $874 \leq \text{FRS}_{1:11} \leq 896$  then
  sign  $\leftarrow$  FPR(FRS)0
  exp  $\leftarrow$  FPR(FRS)1:11 - 1023
  frac  $\leftarrow$  0b1 || FPR(FRS)12:63
  denormalize operand
  do while exp < -126
    frac  $\leftarrow$  0b0 || frac0:62
    exp  $\leftarrow$  exp + 1
  WORD0  $\leftarrow$  sign
  WORD1:8  $\leftarrow$  0x00
  WORD9:31  $\leftarrow$  frac1:23
else WORD  $\leftarrow$  undefined

```

Note that if the value to be stored by a single-precision store floating-point instruction exceeds the maximum number representable in single-precision format, the first case above (no denormalization required) applies. The result stored in WORD is then a well-defined value, but is not numerically equal to the value in the source register (that is, the result of a single-precision load floating-point from WORD does not compare equal to the contents of the original source register).

For double-precision store floating-point instructions and for the Store Floating-Point as Integer Word instruction, no conversion is required, as the data from the FPR is copied directly into memory.

Many floating-point store instructions have an update form, in which GPR(rA) is updated with the EA. For these forms, if rA≠0, the EA is placed into GPR(rA).

Floating-point store accesses cause a data storage interrupt if the program is not allowed to write to the location. Integer store accesses cause a data TLB error interrupt if the program attempts to access memory that is unavailable. Store instructions are shown in [Table 71](#).

Book E supports both big-endian and little-endian byte ordering.

**Table 71. Floating-point store instructions**

Instruction	Mnemonic	Syntax
Store floating-point double	<b>stfd</b>	frS,D(rA)
Store floating-point double with update	<b>stfdu</b>	frS,D(rA)
Store floating-point double extended	<b>stfde</b>	frS,DES(rA)
Store floating-point double with update extended	<b>stfdue</b>	frS,DES(rA)
Store floating-point double indexed	<b>stfdx</b>	frS,rA,rB
Store floating-point double with update indexed	<b>stfdux</b>	frS,rA,rB
Store floating-point double indexed extended	<b>stfdxe</b>	frS,rA,rB
Store floating-point double with update indexed extended	<b>stfduxe</b>	frS,rA,rB
Store floating-point as integer word indexed	<b>stfiwx</b>	frS,rA,rB
Store floating-point as integer word indexed extended	<b>stfiwxe</b>	frS,rA,rB
Store floating-point single	<b>stfs</b>	frS,D(rA)
Store floating-point single with update	<b>stfsu</b>	frS,D(rA)

Table 71. Floating-point store instructions (continued)

Instruction	Mnemonic	Syntax
Store floating-point single extended	<b>stfse</b>	frS,DES(rA)
Store floating-point single with update extended	<b>stfsue</b>	frS,DES(rA)
Store floating-point single indexed	<b>stfsx</b>	frS,rA,rB
Store floating-point single with update indexed	<b>stfsux</b>	frS,rA,rB
Store floating-point single indexed extended	<b>stfsxe</b>	frS,rA,rB
Store floating-point single with update indexed extended	<b>stfsuxe</b>	frS,rA,rB

#### 4.3.1.3.3 Floating-point move instructions

Described in [Table 72](#), these instructions copy data from one FPR to another, altering the sign bit (bit 0) as described below for **fneg**, **fabs**, and **fnabs**. These instructions treat NaNs just like any other kind of value (for example, the sign bit of a NaN may be altered by **fneg**, **fabs**, and **fnabs**). These instructions do not alter the FPSCR.

Table 72. Floating-point move instructions

Instruction	Mnemonic	Syntax
Floating Absolute Value	<b>fabs[.]</b>	frD,frB
Floating Move Register	<b>fmr[.]</b>	frD,frB
Floating Negative Absolute Value	<b>fnabs[.]</b>	frD,frB
Floating Negate	<b>fneg[.]</b>	frD,frB

#### 4.3.1.4 Floating-point arithmetic instructions

The following sections describe elementary arithmetic, multiply-add, rounding/conversion, compare, and status/control instructions.

##### 4.3.1.4.1 Floating-point elementary arithmetic instructions

[Table 73](#) lists mnemonics and syntax of floating-point elementary arithmetic instructions.

Table 73. Floating-point elementary arithmetic instructions

Instruction	Mnemonic	Syntax
Floating add	<b>fadd[.]</b>	frD,frA,frB
Floating add single	<b>fadds[.]</b>	frD,frA,frB
Floating divide	<b>fdiv[.]</b>	frD,frA,frB
Floating divide single	<b>fdivs[.]</b>	frD,frA,frB
Floating multiply	<b>fmul[.]</b>	frD,frA,frC
Floating multiply single	<b>fmuls[.]</b>	frD,frA,frC
Floating reciprocal estimate single	<b>fres[.]</b>	frD,frB
Floating reciprocal square root estimate	<b>frsqrte[.]</b>	frD,frB

**Table 73. Floating-point elementary arithmetic instructions (continued)**

Instruction	Mnemonic	Syntax
Floating square root	<b>fsqrt[.]</b>	frD,frB
Floating square root single	<b>fsqrts[.]</b>	frD,frB
Floating subtract	<b>fsub[.]</b>	frD,frA,frB
Floating subtract single	<b>fsubs[.]</b>	frD,frA,frB

#### 4.3.1.4.2 Floating-point multiply-add instructions

These instructions combine a multiply and an add operation without an intermediate rounding operation. FPSCR status bits, described in [Table 74](#) are set as follows:

- Overflow, underflow, and inexact exception bits, the FR, FI, and FPRF fields are set based on the final result of the operation, not on the result of the multiplication.
- Invalid operation exception bits are set as if the multiplication and the addition were performed using two separate instructions (**fmul[s]**, followed by **fadd[s]** or **fsub[s]**). That is, any of the following actions will cause appropriate exception bits to be set:
  - Multiplication of infinity by 0
  - Multiplication of anything by an SNaN
  - Addition of anything with an SNaN

**Table 74. Floating-point multiply-add instructions**

Instruction	Mnemonic	Instruction
Floating Multiply-Add	<b>fmadd[.]</b>	frD,frA,frB,frC
Floating Multiply-Add Single	<b>fmadds[.]</b>	frD,frA,frB,frC
Floating Multiply-Subtract	<b>fmsub[.]</b>	frD,frA,frB,frC
Floating Multiply-Subtract Single	<b>fmsubs[.]</b>	frD,frA,frB,frC
Floating Negative Multiply-Add	<b>fnmadd[.]</b>	frD,frA,frB,frC
Floating Negative Multiply-Add Single	<b>fnmadds[.]</b>	frD,frA,frB,frC
Floating Negative Multiply-Subtract	<b>fnmsub[.]</b>	frD,frA,frB,frC
Floating Negative Multiply-Subtract Single	<b>fnmsubs[.]</b>	frD,frA,frB,frC

#### 4.3.1.5 Floating-point rounding and conversion instructions

**Table 75. Floating-point rounding and conversion instructions**

Instruction	Mnemonic	Syntax
Floating Convert from Integer Double Word	<b>fcfid</b>	frD,frB
Floating Convert to Integer Double Word	<b>fctid</b>	frD,frB
Floating Convert to Integer Double word and round to Zero	<b>fctidz</b>	frD,frB
Floating Convert to Integer Word	<b>fctiw[.]</b>	frD,frB

**Table 75. Floating-point rounding and conversion instructions (continued)**

Instruction	Mnemonic	Syntax
Floating Convert to Integer Word and Round to Zero	<b>fctiwz</b> [.]	<b>frD,frB</b>
Floating Round to Single-Precision	<b>frsp</b> [.]	<b>frD,frB</b>

#### 4.3.1.6 Floating-point compare instructions

The floating-point compare instructions compare the contents of two FPRs. Comparison ignores the sign of zero (that is, regards +0 as equal to -0). The comparison result can be ordered or unordered. The comparison sets one bit in the designated CR field and clears the other three. The floating-point condition code, FPSCR[FPCC], is set in the same way.

The CR field and the FPCC are set as described in [Table 76](#).

**Table 76. CR field settings**

Bit	Name	Description
0	FL	$(frA) < (frB)$
1	FG	$(frA) > (frB)$
2	FE	$(frA) = (frB)$
3	FU	$(frA) ? (frB)$ (unordered)

The floating-point compare and select instruction set is shown in [Table 77](#).

**Table 77. Floating-point compare and select instructions**

Instruction	Mnemonic	Syntax
Floating Compare Ordered	<b>fcmpo</b>	<b>crD,frA,frB</b>
Floating Compare Unordered	<b>fcmpu</b>	<b>crD,frA,frB</b>
Floating Select	<b>fsel</b> <b>fsel.</b>	<b>frD,frA,frB,frC</b> <b>frD,frA,frB,frC</b>

#### 4.3.1.7 Floating-point status and control register instructions

Every FPSCR instruction synchronizes the effects of all floating-point instructions executed by a given processor. Executing a FPSCR instruction ensures that all floating-point instructions previously initiated by the given processor have completed before the FPSCR instruction is initiated, and that no subsequent floating-point instructions are initiated by the given processor until the FPSCR instruction completes. In particular:

- All exceptions caused by the previously initiated instructions are recorded in the FPSCR before the FPSCR instruction is initiated.
- All invocations of floating-point enabled exception-type program interrupt that will be caused by the previously initiated instructions have occurred before the FPSCR instruction is initiated.
- No subsequent floating-point instruction that depends on or alters the settings of any FPSCR bits is initiated until the FPSCR instruction has completed.

Floating-point load and floating-point store instructions ([Table 78](#)) are not affected.

Table 78. Floating-point status and control register instructions

Instruction	Mnemonic	Syntax
Move from FPSCR	<b>mffs</b> <b>mffs.</b>	<b>frD</b> <b>frD</b>
Move to FPSCR Bit 0	<b>mtfsb0</b> <b>mtfsb0.</b>	<b>crbD</b> <b>crbD</b>
Move to FPSCR Bit 1	<b>mtfsb1</b> <b>mtfsb1.</b>	<b>crbD</b> <b>crbD</b>
Move to FPSCR Fields	<b>mtfsf</b> <b>mtfsf.</b>	FM,frB FM,frB
Move to FPSCR Field Immediate	<b>mtfsfi</b> <b>mtfsfi.</b>	crD,IMM crD,IMM

### 4.3.1.8 Load and store instructions

Load and store instructions are issued and translated in program order; however, the accesses can occur out of order. Synchronizing instructions are provided to enforce strict ordering. The following load and store instructions are defined:

- Integer load instructions
- Integer store instructions
- Integer load and store with byte-reverse instructions
- Integer load and store multiple instructions
- Memory synchronization instructions
- SPE APU load and store instructions for reading and writing 64-bit GPRs. Some of these instructions are also implemented by processors that support the embedded vector single-precision and embedded scalar double-precision floating-point APUs, which use the extended 64-bit GPRs. See [Section 4.6.1: SPE and embedded floating-point APUs](#).

#### 4.3.1.8.1 Self-modifying code

When a processor modifies any memory location that can contain an instruction, **software must ensure that** the instruction cache is made consistent with data memory and that the modifications are made visible to the instruction fetching mechanism. This must be done even if the cache is disabled or if the page is marked caching-inhibited.

The following instruction sequence can be used to accomplish this when the instructions being modified are in memory that is memory-coherence required and one processor both modifies the instructions and executes them. (Additional synchronization is needed when one processor modifies instructions that another processor will execute.)

The following sequence synchronizes the instruction stream (using either **dcbst** or **dcbf**):

```

dcbst (or dcbf) | update memory
msync | wait for update
icbi | remove (invalidate) copy in instruction cache
msync | ensure the ICBI invalidate is complete
isync | remove copy in own instruction buffer

```

These operations are required because the data cache is a write-back cache. Because instruction fetching bypasses the data cache, changes to items in the data cache cannot be reflected in memory until the fetch operations complete. The **msync** after the **icbi** is required to ensure that the **icbi** invalidation has completed in the instruction cache.

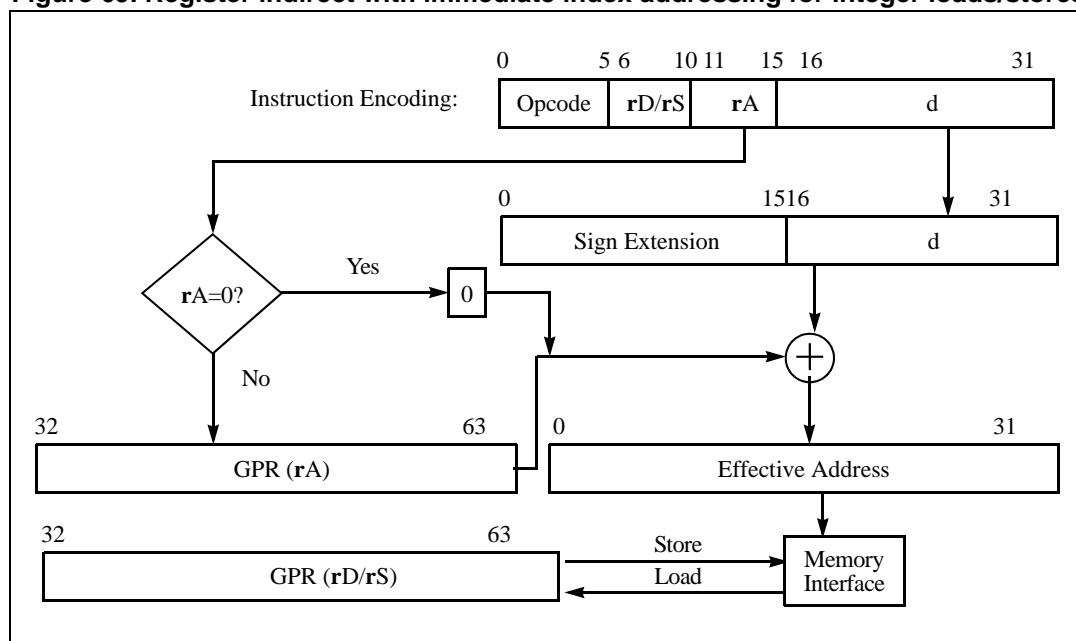
Special care must be taken to avoid coherency paradoxes in systems that implement unified secondary caches, and designers should carefully follow the guidelines for maintaining cache coherency discussed in the user’s manual.

Integer load and store address generation

Integer load and store operations generate EAs using register indirect with immediate index mode, register indirect with index mode, or register indirect mode, which are described as follows:

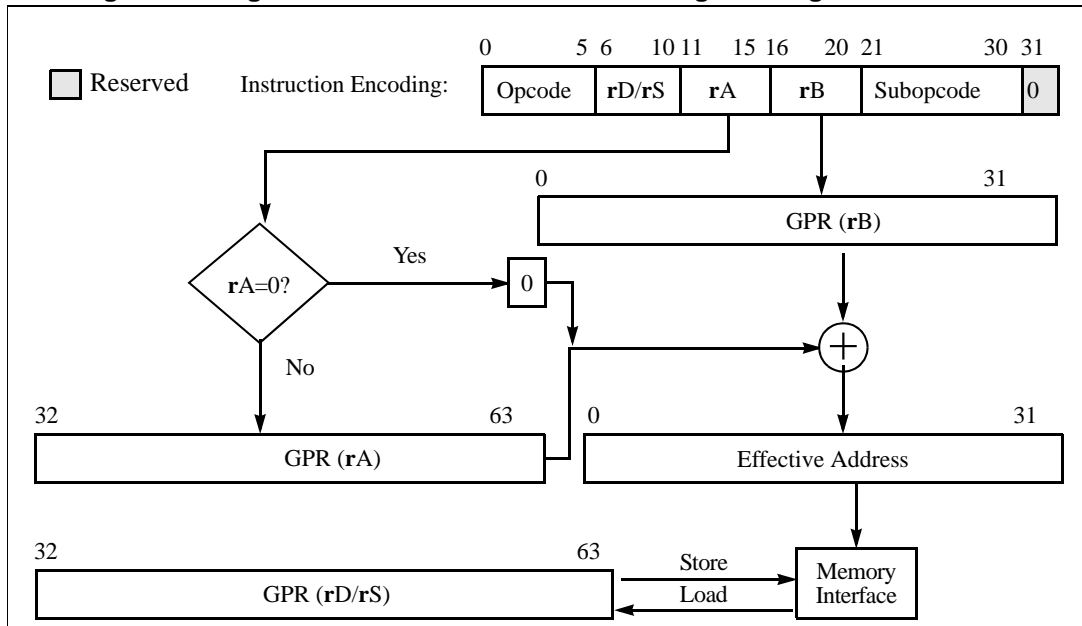
- Register indirect with immediate index addressing for integer loads and stores. Instructions using this addressing mode contain a signed 16-bit immediate index (d operand), which is sign extended and added to the contents of a general-purpose register specified in the instruction (rA operand), to generate the EA. If r0 is specified, a value of zero is added to the immediate index (d operand) in place of the contents of r0. The option to specify rA or 0 is shown in the instruction descriptions as (rA|0). *Figure 69* shows how an EA is generated using this mode.

**Figure 69. Register indirect with immediate index addressing for integer loads/stores**



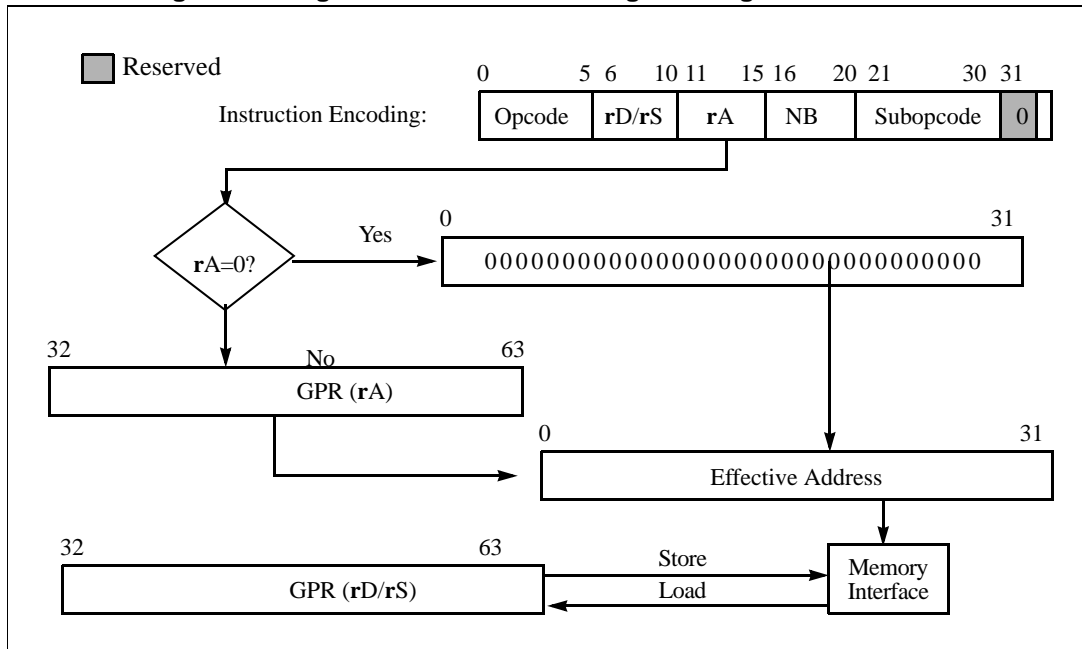
- Register indirect with index addressing for integer loads and stores. Instructions using this mode cause the contents of two GPRs (specified as operands rA and rB) to be added in the EA generation. A zero in place of the rA operand causes a zero to be added to the GPR contents specified in operand rB. The option to specify rA or 0 is shown in the instruction descriptions as (rA|0). *Figure 70* shows how an EA is generated using this mode.

Figure 70. Register indirect with index addressing for integer loads/stores



- Register indirect addressing for integer loads and stores. Instructions using this addressing mode use the contents of the GPR specified by the rA operand as the EA. A zero in the rA operand generates an EA of zero. The option to specify rA or 0 is shown in the instruction descriptions as (rA|0). Figure 71 shows how an EA is generated using this mode.

Figure 71. Register indirect addressing for integer loads/stores



See [Section 4.2.3.3: Effective address calculation](#), for information about calculating EAs. Note that in some implementations, operations that are not naturally aligned can suffer performance degradation. [Section 5.7.6: Alignment interrupt](#), for additional information about load and store address alignment interrupts.



### 4.3.1.9 Register indirect integer load instructions

For integer load instructions, the byte, half word, or word addressed by the EA is loaded into rD. Many integer load instructions have an update form, in which rA is updated with the generated EA. For these forms, if  $rA \neq 0$  and  $rA \neq rD$  (otherwise invalid), the EA is placed into rA and the memory element (byte, half word, or word) addressed by the EA is loaded into rD. Note that the Book E architecture defines load with update instructions with operand  $rA = 0$  or  $rA = rD$  as invalid forms.

### 4.3.1.10 Integer load instructions

Table 79. Integer load instructions

Name	Mnemonic	Syntax
Load Byte and Zero	lbz	rD,d(rA)
Load Byte and Zero Indexed	lbzx	rD,rA,rB
Load Byte and Zero with Update	lbzu	rD,d(rA)
Load Byte and Zero with Update Indexed	lbzux	rD,rA,rB
Load Half Word and Zero	lhz	rD,d(rA)
Load Half Word and Zero Indexed	lhzx	rD,rA,rB
Load Half Word and Zero with Update	lhzu	rD,d(rA)
Load Half Word and Zero with Update Indexed	lhzux	rD,rA,rB
Load Half Word Algebraic	lha	rD,d(rA)
Load Half Word Algebraic Indexed	lhax	rD,rA,rB
Load Half Word Algebraic with Update	lhau	rD,d(rA)
Load Half Word Algebraic with Update Indexed	lhaux	rD,rA,rB
Load Word and Zero	lwz	rD,d(rA)
Load Word and Zero Indexed	lwzx	rD,rA,rB
Load Word and Zero with Update	lwzu	rD,d(rA)
Load Word and Zero with Update Indexed	lwzux	rD,rA,rB

#### 4.3.1.10.1 Integer store instructions

For integer store instructions, the rS contents are stored into the byte, half word, word or double word in memory addressed by the EA. Many store instructions have an update form in which rA is updated with the EA. For these forms, the following rules apply:

- If  $rA \neq 0$ , the EA is placed into rA.
- If  $rS = rA$ , the contents of register rS are copied to the target memory element and the generated EA is placed into rA (rS).

The Book E architecture defines store with update instructions with  $rA = 0$  as an invalid form. In addition, it defines integer store instructions with the CR update option enabled (Rc field, bit 31, in the instruction encoding = 1) to be an invalid form. [Table 80](#) summarizes integer store instructions.

Table 80. Integer store instructions

Name	Mnemonic	Syntax
Store Byte	stb	rS,d(rA)
Store Byte Indexed	stbx	rS,rA,rB
Store Byte with Update	stbu	rS,d(rA)
Store Byte with Update Indexed	stbux	rS,rA,rB
Store Half Word	sth	rS,d(rA)
Store Half Word Indexed	sthx	rS,rA,rB
Store Half Word with Update	sthv	rS,d(rA)
Store Half Word with Update Indexed	sthvx	rS,rA,rB
Store Word	stw	rS,d(rA)
Store Word Indexed	stwx	rS,rA,rB
Store Word with Update	stwu	rS,d(rA)
Store Word with Update Indexed	stwux	rS,rA,rB

#### 4.3.1.10.2 Integer load and store with byte-reverse instructions

[Table 81](#) describes integer load and store with byte-reverse instructions. These instructions were defined in part to support the original PowerPC definition of little-endian byte ordering. Note that Book E supports true little endian on a per-page basis. For more information, see [Section 4.2.3.4: Byte ordering](#).

Table 81. Integer load and store with byte-reverse instructions

Name	Mnemonic	Syntax
Load Half Word Byte-Reverse Indexed	lhbrx	rD,rA,rB
Load Word Byte-Reverse Indexed	lwbrx	rD,rA,rB
Store Half Word Byte-Reverse Indexed	sthbrx	rS,rA,rB
Store Word Byte-Reverse Indexed	stwbrx	rS,rA,rB

#### 4.3.1.10.3 Integer load and store multiple instructions

The load/store multiple instructions are used to move blocks of data to and from the GPRs. The load multiple and store multiple instructions can have operands that require memory accesses crossing a 4-Kbyte page boundary. As a result, these instructions can be interrupted by a data storage interrupt associated with the address translation of the second page.

*Note:* If one of these instructions is interrupted, it may be restarted, requiring multiple memory accesses.

The Book E architecture defines the Load Multiple Word (**lmw**) instruction ([Table 82](#)) with rA in the range of registers to be loaded as an invalid form. Load and store multiple accesses must be word aligned; otherwise, they cause an alignment exception.

**Table 82. Integer load and store multiple instructions**

Name	Mnemonic	Syntax
Load Multiple Word	lmw	rD,d(rA)
Store Multiple Word	stmw	rS,d(rA)

#### 4.3.1.10.4 Integer load and store string instructions

The integer load and store string instructions allow movement of data from memory to registers or from registers to memory without concern for alignment. These instructions can be used for a short move between arbitrary memory locations or to initiate a long move between misaligned memory fields. However, in some implementations, these instructions are likely to have greater latency and take longer to execute, perhaps much longer, than a sequence of individual load or store instructions that produce the same results.

[Table 83](#) summarizes the integer load and store string instructions.

**Table 83. Integer load and store string instructions**

Name	Mnemonic	Syntax
Load String Word Immediate	lswi	rD,rA,NB
Load String Word Indexed	lswx	rD,rA,rB
Store String Word Immediate	stswi	rS,rA,NB
Store String Word Indexed	stswx	rS,rA,rB

Load string and store string instructions can involve operands that are not word-aligned.

#### 4.3.1.10.5 Floating-point load and store address generation

Floating-point load and store operations, listed in [Table 84](#), generate EAs using the register indirect with immediate index addressing mode and register indirect with index addressing mode. Floating-point loads and stores are not supported for direct-store accesses. The use of floating-point loads and stores for direct-store accesses results in an alignment interrupt.

There are two forms of the floating-point load instruction—single-precision and double-precision operand formats. Because the FPRs support only the floating-point double-precision format, single-precision floating-point load instructions convert single-precision data to double-precision format before loading an operand into an FPR.

The floating-point load and store indexed instructions (**lfsx**, **lfsux**, **lfdx**, **lfdx**, **stfsx**, **stfsux**, **stfdx**, and **stfdx**) are invalid when the Rc bit is one.

The PowerPC architecture defines load with update with rA = 0 as an invalid form.

**Table 84. Floating-point load instructions**

Name	Mnemonic	Syntax
Load Floating-Point Single	lfs	frD,d(rA)
Load Floating-Point Single Indexed	lfsx	frD,rA,rB
Load Floating-Point Single with Update	lfsu	frD,d(rA)

**Table 84. Floating-point load instructions (continued)**

Name	Mnemonic	Syntax
Load Floating-Point Single with Update Indexed	lfsux	frD,rA,rB
Load Floating-Point Double	lfd	frD,d(rA)
Load Floating-Point Double Indexed	lfdx	frD,rA,rB
Load Floating-Point Double with Update	lfdv	frD,d(rA)
Load Floating-Point Double with Update Indexed	lfdvx	frD,rA,rB

**4.3.1.10.6 Floating-point store instructions**

This section describes floating-point store instructions. There are three basic forms of the store instruction—single-precision, double-precision, and integer. The integer form is supported by the optional **stfiwx** instruction. Because the FPRs support only double-precision format for floating-point data, single-precision floating-point store instructions convert double-precision data to single-precision format before storing the operands. [Table 85](#) summarizes the floating-point store instructions.

**Table 85. Floating-point store instructions**

Name	Mnemonic	Syntax
Store Floating-Point Single	stfs	frS,d(rA)
Store Floating-Point Single Indexed	stfsx	frS,r B
Store Floating-Point Single with Update	stfsu	frS,d(rA)
Store Floating-Point Single with Update Indexed	stfsux	frS,r B
Store Floating-Point Double	stfd	frS,d(rA)
Store Floating-Point Double Indexed	stfdx	frS,rB
Store Floating-Point Double with Update	stfdv	frS,d(rA)
Store Floating-Point Double with Update Indexed	stfdvx	frS,rB
Store Floating-Point as Integer Word Indexed <sup>(1)</sup>	stfiwx	frS,rB

1. The **stfiwx** instruction is optional to the Book E architecture.

Some floating-point store instructions require conversions in the LSU. [Table 86](#) shows conversions the LSU makes when executing a Store Floating-Point Single instruction.

**Table 86. Store floating-point single behavior**

FPR Precision	Data Type	Action
Single	Normalized	Store
Single	Denormalized	Store
Single	Zero, infinity, QNaN	Store
Single	SNaN	Store



**Table 86. Store floating-point single behavior (continued)**

FPR Precision	Data Type	Action
Double	Normalized	If ( $\text{exp} \leq 896$ ) then denormalize and store, else store
Double	Denormalized	Store zero
Double	Zero, infinity, QNaN	Store
Double	SNaN	Store

[Table 87](#) shows the conversions made when performing a Store Floating-Point Double instruction. Most entries in the table indicate that the floating-point value is simply stored. Only in a few cases are any other actions taken.

**Table 87. Store floating-point double behavior**

FPR Precision	Data Type	Action
Single	Normalized	Store
Single	Denormalized	Normalize and store
Single	Zero, infinity, QNaN	Store
Single	SNaN	Store
Double	Normalized	Store
Double	Denormalized	Store
Double	Zero, infinity, QNaN	Store
Double	SNaN	Store

#### 4.3.1.11 Branch and flow control instructions

Some branch instructions can redirect instruction execution conditionally based on the value of bits in the CR.

##### 4.3.1.11.1 Branch instruction address calculation

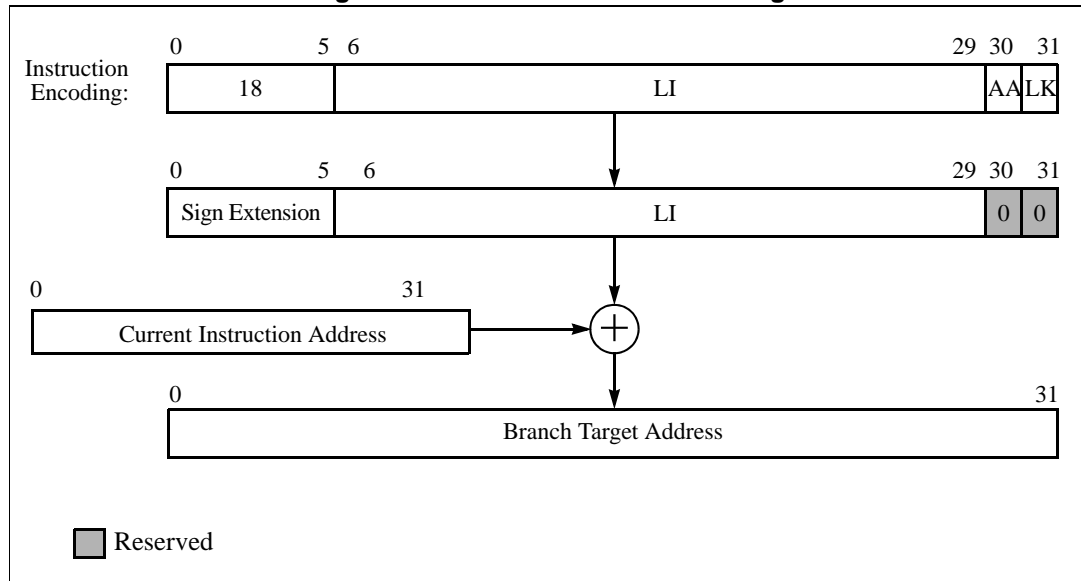
Branch instructions can alter the sequence of instruction execution. Instruction addresses are always assumed to be word aligned; the Book E processors ignore the two low-order bits of the generated branch target address. Branch instructions compute the EA of the next instruction address using the following addressing modes:

- Branch relative
- Branch conditional to relative address
- Branch to absolute address
- Branch conditional to absolute address
- Branch conditional to link register (LR)
- Branch conditional to count register (CTR)

### 4.3.1.11.2 Branch relative addressing mode

Instructions that use branch relative addressing generate the next instruction address by sign extending and appending 0b00 to the immediate displacement operand **LI**, and adding the resultant value to the current instruction address. Branches using this mode have the absolute addressing option disabled (AA field, bit 30, in the instruction encoding = 0). The LR update option can be enabled (LK field, bit 31, in the instruction encoding = 1). This causes the EA of the instruction following the branch instruction to be placed in the LR. [Figure 72](#) shows how the branch target address is generated using this mode.

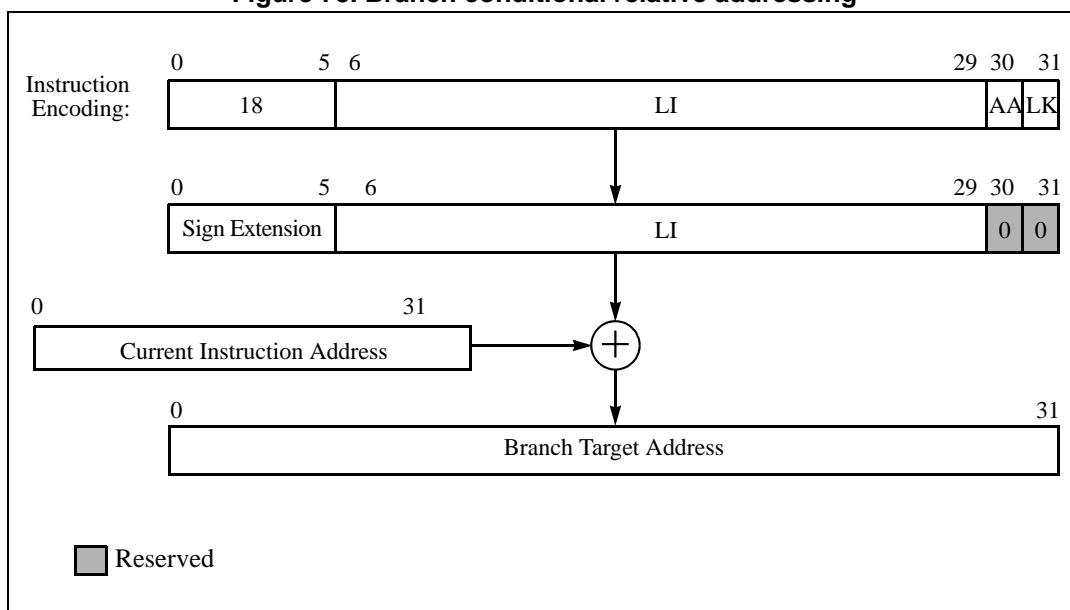
**Figure 72. Branch relative addressing**



### 4.3.1.11.3 Branch conditional to relative addressing mode

If branch conditions are met, instructions that use the branch conditional to relative addressing mode generate the next instruction address by sign extending and appending results to the immediate displacement operand (**BD**) and adding the resultant value to the current instruction address. Branches using this mode have the absolute addressing option disabled (AA field, bit 30, in the instruction encoding = 0). The LR update option can be enabled (LK field, bit 31, in the instruction encoding = 1). This option causes the EA of the instruction following the branch instruction to be placed in the LR. [Figure 73](#) shows how the branch target address is generated using this mode.

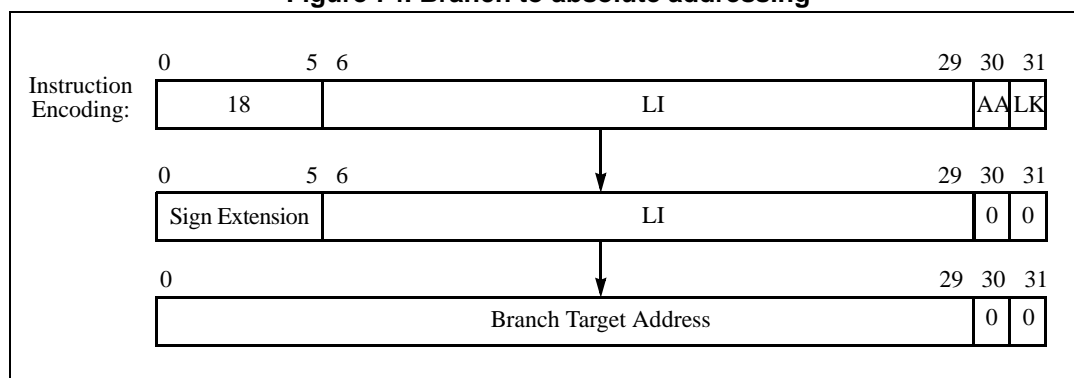
Figure 73. Branch conditional relative addressing



4.3.1.11.4 Branch to absolute addressing mode

Instructions that use branch to absolute addressing mode generate the next instruction address by sign extending and appending 0b00 to the LI operand. Branches using this addressing mode have the absolute addressing option enabled (AA field, bit 30, in the instruction encoding = 1). The LR update option can be enabled (LK field, bit 31, in the instruction encoding = 1). This option causes the EA of the instruction following the branch instruction to be placed in the LR. *Figure 74* shows how the branch target address is generated using this mode.

Figure 74. Branch to absolute addressing

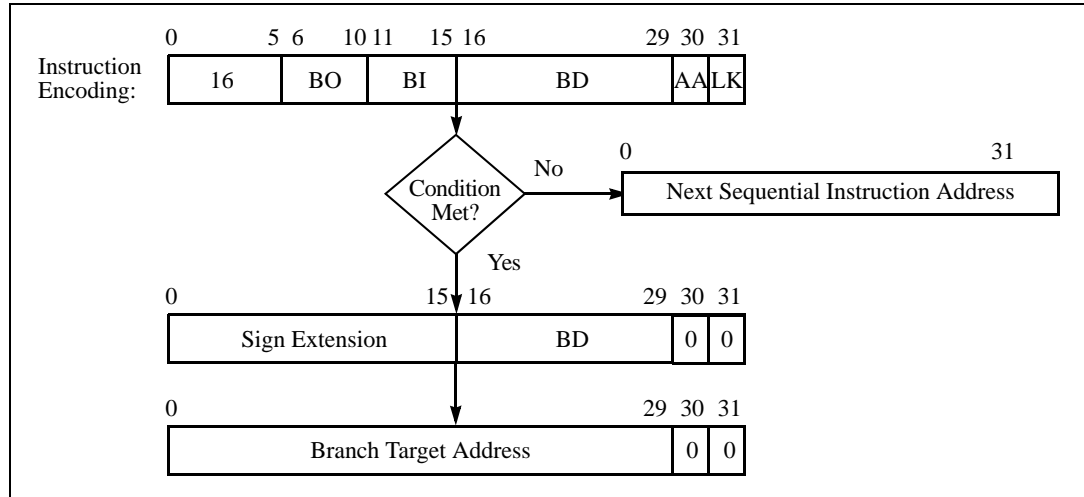


4.3.1.11.5 Branch conditional to absolute addressing mode

If the branch conditions are met, instructions that use the branch conditional to absolute addressing mode generate the next instruction address by sign extending and appending 0b00 to the BD operand. Branches using this addressing mode have the absolute addressing option enabled (AA field, bit 30, in the instruction encoding = 1). The LR update option can be enabled (bit 31 (LK) in the instruction encoding = 1). This option causes the

EA of the instruction following the branch instruction to be placed in the LR. *Figure 75* shows how the branch target address is generated using this mode.

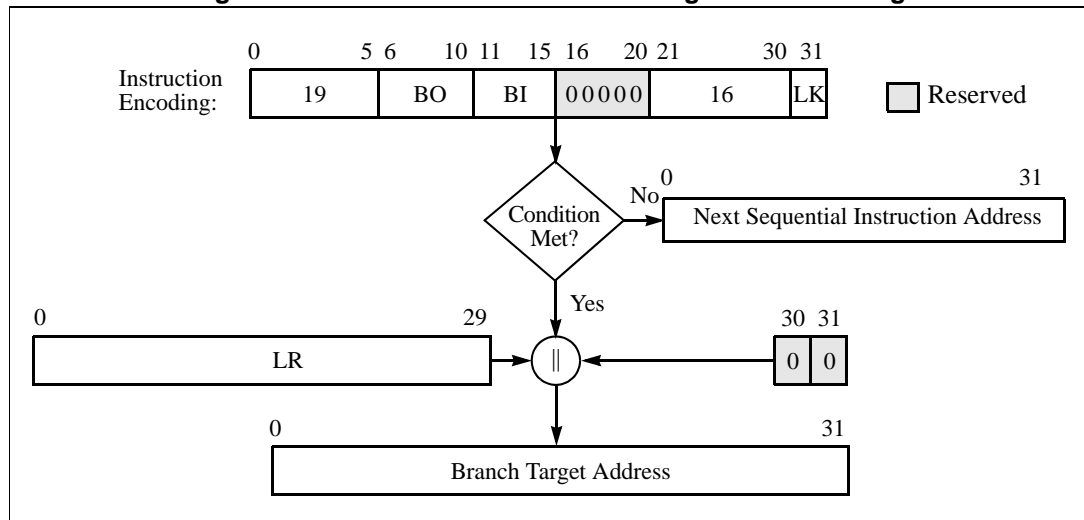
**Figure 75. Branch conditional to absolute addressing**



**4.3.1.11.6 Branch conditional to link register addressing mode**

If the branch conditions are met, the branch conditional to LR instruction generates the next instruction address by fetching the contents of the LR and clearing the two low-order bits to zero. The LR update option can be enabled (LK field, bit 31, in the instruction encoding = 1). This option causes the EA of the instruction following the branch instruction to be placed in the LR. *Figure 76* shows how the branch target address is generated using this mode.

**Figure 76. Branch conditional to link register addressing**



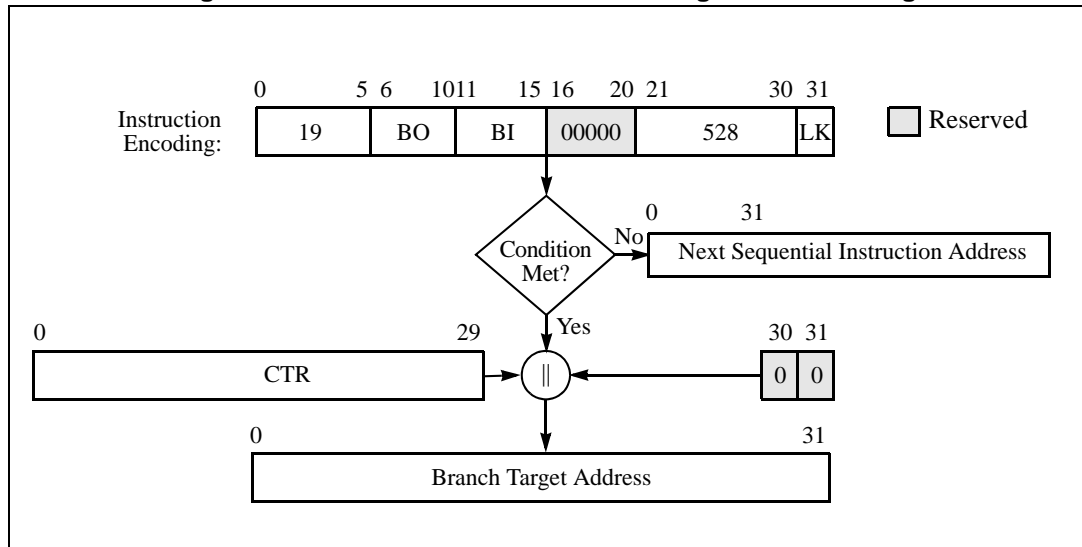
**4.3.1.11.7 Branch conditional to count register addressing mode**

If the branch conditions are met, the branch conditional to count register instruction generates the next instruction address by fetching the contents of the count register (CTR)



and clearing the two low-order bits to zero. The LR update option can be enabled (LK field, bit 31, in the instruction encoding = 1). This option causes the EA of the instruction following the branch instruction to be placed in the LR. [Figure 77](#) shows how the branch target address is generated when using this mode.

**Figure 77. Branch conditional to count register addressing**



**4.3.1.12 Conditional branch control**

*Note:* Some processors do not implement the static branch prediction defined in Book E and described here. For those processors, the BO operand is ignored for branch prediction.

For branch conditional instructions, the BO operand specifies the conditions under which the branch is taken. The first four bits of the BO operand specify how the branch is affected by or affects the condition and count registers. The fifth bit, shown in [Table 89](#) as having the value *y*, is used by some implementations for branch prediction as described below.

**Table 88. BO bit descriptions**

BO Bits	Description
0	Setting this bit causes the CR bit to be ignored.
1	Bit value to test against
2	Setting this causes the decrement to not be decremented.
3	Setting this bit reverses the sense of the CTR test.
4	Used for the <i>y</i> bit, which provides a hint about whether a conditional branch is likely to be taken and may be used by some implementations to improve performance.

The encodings for the BO operands are shown in [Table 89](#).

**Table 89. BO operand encodings<sup>(1)</sup>**

BO	Description
0000y	Decrement the CTR, then branch if the decremented CTR ≠ 0 and the condition is FALSE.
0001y	Decrement the CTR, then branch if the decremented CTR = 0 and the condition is FALSE.
001zy	Branch if the condition is FALSE.
0100y	Decrement the CTR, then branch if the decremented CTR ≠ 0 and the condition is TRUE.
0101y	Decrement the CTR, then branch if the decremented CTR = 0 and the condition is TRUE.
011zy	Branch if the condition is TRUE.
1z00y	Decrement the CTR, then branch if the decremented CTR ≠ 0.
1z01y	Decrement the CTR, then branch if the decremented CTR = 0.
1z1zz	Branch always.

1. In this table, z indicates a bit that is ignored. Note that the z bits should be cleared, as they may be assigned a meaning in some future version of the architecture. The y bit provides a hint about whether a conditional branch is likely to be taken and may be used by some implementations to improve performance.

The branch always encoding of the BO operand does not have a y bit.

Clearing the y bit indicates a predicted behavior for the branch instruction as follows:

- For **bcx** with a negative value in the displacement operand, the branch is taken.
- In all other cases (**bcx** with a non-negative value in the displacement operand, **bclrx**, or **bcctrx**), the branch is not taken.

Setting the y bit reverses the preceding indications.

The sign of the displacement operand is used as described above even if the target is an absolute address. The default value for the y bit should be 0 and should be set to 1 only if software has determined that the prediction corresponding to y = 1 is more likely to be correct than the prediction corresponding to y = 0. Software that does not compute branch predictions should clear the y bit.

In most cases, the branch should be predicted to be taken if the value of the following expression is 1, and predicted to fall through if the value is 0.

$$((BO[0] \& BO[2]) | S) \approx BO[4]$$

In the expression above, S (bit 16 of the branch conditional instruction coding) is the sign bit of the displacement operand if the instruction has a displacement operand and is 0 if the operand is reserved. BO[4] is the y bit, or 0 for the branch always encoding of the BO operand. (Advantage is taken of the fact that, for **bclrx** and **bcctrx**, bit 16 of the instruction is part of a reserved operand and therefore must be 0.)

The 5-bit BI operand in branch conditional instructions specifies which CR bit represents the condition to test. The CR bit selected is BI +32, as shown in [Table 17](#).

If the branch instructions contain immediate addressing operands, the target addresses can be computed sufficiently ahead of the branch instruction that instructions can be fetched along the target path. If the branch instructions use the link and count registers, instructions along the target path can be fetched if the link or count register is loaded sufficiently ahead of the branch instruction.

Branching can be conditional or unconditional, and optionally a branch return address is created by storing the EA of the instruction following the branch instruction in the LR after the branch target address has been computed. This is done regardless of whether the branch is taken.

#### 4.3.1.12.1 Branch instructions

[Table 90](#) lists branch instructions provided by the Book E processors. A set of simplified mnemonics and symbols is provided for the most frequently used forms of branch conditional, compare, trap, rotate and shift, and certain other instructions. See [Appendix B: Simplified mnemonics for PowerPC instructions](#).

**Table 90. Branch instructions**

Name	Mnemonic	Syntax
Branch	<b>b</b> ( <b>ba bl bla</b> )	target_addr
Branch Conditional	<b>bc</b> ( <b>bca bcl bcla</b> )	BO,BI,target_addr
Branch Conditional to Link Register	<b>bclr</b> ( <b>bclrl</b> )	BO,BI
Branch Conditional to Count Register	<b>bcctr</b> ( <b>bcctrl</b> )	BO,BI

Note that the EIS defines the Integer Select instruction, **isel**, which can be used to more efficiently handle sequences with multiple conditional branches. Its syntax is given in [Section 4.6.2](#). A detailed description including an example of how **isel** can be used can be found in [Section 8.1.2](#).

#### 4.3.1.12.2 Condition register (cr) logical instructions

CR logical instructions, shown in [Table 91](#), and the Move Condition Register Field (**mcrf**) instruction are also defined as flow control instructions.

**Table 91. Condition register logical instructions**

Name	Mnemonic	Syntax
Condition Register AND	<b>crand</b>	<b>crbD,crbA,crbB</b>
Condition Register OR	<b>cror</b>	<b>crbD,crbA,crbB</b>
Condition Register XOR	<b>crxor</b>	<b>crbD,crbA,crbB</b>
Condition Register NAND	<b>crnand</b>	<b>crbD,crbA,crbB</b>
Condition Register NOR	<b>crnor</b>	<b>crbD,crbA,crbB</b>
Condition Register Equivalent	<b>creqv</b>	<b>crbD,crbA,crbB</b>
Condition Register AND with Complement	<b>crandc</b>	<b>crbD,crbA,crbB</b>
Condition Register OR with Complement	<b>crorc</b>	<b>crbD,crbA,crbB</b>
Move Condition Register Field	<b>mcrf</b>	<b>crfD,crfS</b>

Note that if the LR update option is enabled for any of these instructions, the Book E architecture defines these forms of the instructions as invalid.

**4.3.1.12.3 Trap instructions**

The trap instructions shown in [Table 92](#) test for a specified set of conditions. If any of the conditions tested by a trap instruction are met, the system trap type program interrupt is taken. For more information, see [Section 5.7.7: Program interrupt](#). If the tested conditions are not met, instruction execution continues normally. See [Appendix B: Simplified mnemonics for PowerPC instructions](#).

**Table 92. Trap instructions**

Name	Mnemonic	Syntax
Trap Word Immediate	<b>twi</b>	TO,rA,SIMM
Trap Word	<b>tw</b>	TO,rA,rB

**4.3.1.13 System linkage instruction**

The system call (**sc**) instruction permits a program to call on the system to perform a service; see [Table 93](#) and [Section 4.3.2.1: System linkage instructions](#).

**Table 93. System linkage instruction**

Name	Mnemonic	Syntax
System Call	<b>sc</b>	—

Executing this instruction causes the system call interrupt handler to be invoked. For more information, see [Section 5.7.9](#).

**4.3.1.14 Processor control instructions**

Processor control instructions are used to read from and write to the CR, machine state register (MSR), and special-purpose registers (SPRs).

**4.3.1.14.1 Move to/from condition register instructions**

[Table 94](#) summarizes the instructions for reading from or writing to the CR.

**Table 94. Move to/from condition register instructions**

Name	Mnemonic	Syntax
Move to Condition Register Fields	<b>mtcrf</b>	CRM,rS
Move to Condition Register from XER	<b>mcrxr</b>	crD
Move from Condition Register	<b>mfcrr</b>	rD

**4.3.1.14.2 Move to/from special-purpose register instructions**

[Table 95](#) lists the **mtspr** and **mfspr** instructions.

**Table 95. Move to/from special-purpose register instructions**

Name	Mnemonic	Syntax
Move to Special-Purpose Register	<b>mtspr</b>	SPR,rS
Move from Special-Purpose Register	<b>mfspir</b>	rD,SPR

[Table 96](#) summarizes all SPRs defined in Book E, indicating which are user-level access. The SPR number column lists register numbers used in the instruction mnemonics.

**Table 96. Book E special-purpose registers (by SPR abbreviation)**

SPR	Name	Defined SPR number		Access	Supervisor only	Section/ page
		Decimal	Binary			
CSRR0	Critical save/restore register 0	58	00001 11010	Read/Write	Yes	<a href="#">Section 3.9.1.3 on page 84</a>
CSRR1	Critical save/restore register 1	59	00001 11011	Read/Write	Yes	<a href="#">Section 3.9.1.4 on page 84</a>
CTR	Count register	9	00000 01001	Read/Write	No	<a href="#">Section 3.5.3.1 on page 69</a>
DAC1	Data address compare 1	316	01001 11100	Read/Write	Yes	<a href="#">Section 3.13.4 on page 119</a>
DAC2	Data address compare 2	317	01001 11101	Read/Write	Yes	<a href="#">Section 3.13.4 on page 119</a>
DBCR0	Debug control register 0	308	01001 10100	Read/Write	Yes	<a href="#">Section 3.13.1.1 on page 110</a>
DBCR1	Debug control register 1	309	01001 10101	Read/Write	Yes	<a href="#">Section 3.13.1.2 on page 112</a>
DBCR2	Debug control register 2	310	01001 10110	Read/Write	Yes	<a href="#">Section 3.13.1.3 on page 114</a>
DBSR	Debug status register	304	01001 10000	Read/Clear <sup>(1)</sup>	Yes	<a href="#">Section 3.13.2 on page 117</a>
DEAR	Data exception address register	61	00001 11101	Read/Write	Yes	<a href="#">Section 3.9.1.5 on page 85</a>
DEC	Decrementer	22	00000 10110	Read/Write	Yes	<a href="#">Section 3.8.4 on page 82</a>
DECAR	Decrementer auto-reload	54	00001 10110	Write-only	Yes	<a href="#">Section 3.8.5 on page 83</a>
DVC1	Data value compare 1	318	01001 11110	Read/Write	Yes	<a href="#">Section 3.13.5 on page 119</a>
DVC2	Data value compare 2	319	01001 11111			
ESR	Exception syndrome register	62	00001 11110	Read/Write	Yes	<a href="#">Section 3.9.1.8 on page 86</a>
IAC1	Instruction address compare 1	312	01001 11000	Read/Write	Yes	<a href="#">Section 3.13.3 on page 118</a>

**Table 96. Book E special-purpose registers (by SPR abbreviation) (continued)**

SPR	Name	Defined SPR number		Access	Supervisor only	Section/ page
		Decimal	Binary			
IAC2	Instruction address compare 2	313	01001 11001	Read/Write	Yes	<a href="#">Section 3.13.3 on page 118</a>
IAC3	Instruction address compare 3	314	01001 11010	Read/Write	Yes	<a href="#">Section 3.13.3 on page 118</a>
IAC4	Instruction address compare 4	315	01001 11011	Read/Write	Yes	<a href="#">Section 3.13.3 on page 118</a>
IVOR0	Critical input	400	01100 10000	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR1	Critical input interrupt offset	401	01100 10001	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR2	Data storage interrupt offset	402	01100 10010	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR3	Instruction storage interrupt offset	403	01100 10011	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR4	External input interrupt offset	404	01100 10100	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR5	Alignment interrupt offset	405	01100 10101	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR6	Program interrupt offset	406	01100 10110	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR7	Floating-point unavailable interrupt offset	407	01100 10111	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR8	System call interrupt offset	408	01100 11000	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR9	Auxiliary processor unavailable interrupt offset	409	01100 11001	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR10	Decrementer interrupt offset	410	01100 11010	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR11	Fixed-interval timer interrupt offset	411	01100 11011	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR12	Watchdog timer interrupt offset	412	01100 11100	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR13	Data TLB error interrupt offset	413	01100 11101	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR14	Instruction TLB error interrupt offset	414	01100 11110	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR15	Debug interrupt offset	415	01100 11111	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVPR	Interrupt vector	63	00001 11111	Read/Write	Yes	<a href="#">Section 3.13.3 on page 118</a>

Table 96. Book E special-purpose registers (by SPR abbreviation) (continued)

SPR	Name	Defined SPR number		Access	Supervisor only	Section/ page
		Decimal	Binary			
LR	Link register	8	00000 01000	Read/Write	No	<a href="#">Section 3.5.2 on page 68</a>
PID	Process ID register <sup>(2)</sup>	48	00001 10000	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
PIR	Processor ID register	286	01000 11110	Read only	Yes	<a href="#">Section 3.7.3 on page 76</a>
PVR	Processor version register	287	01000 11111	Read only	Yes	<a href="#">Section 3.7.4 on page 76</a>
SPRG0	SPR general 0	272	01000 10000	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG1	SPR general 1	273	01000 10001	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG2	SPR general 2	274	01000 10010	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG3	SPR general 3	259	01000 00011	Read only	No <sup>(3)</sup>	<a href="#">Section 3.10 on page 92</a>
		275	01000 10011	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG4	SPR general 4	260	01000 00100	Read only	No	<a href="#">Section 3.10 on page 92</a>
		276	01000 10100	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG5	SPR general 5	261	01000 00101	Read only	No	<a href="#">Section 3.10 on page 92</a>
		277	01000 10101	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG6	SPR general 6	262	01000 00110	Read only	No	<a href="#">Section 3.10 on page 92</a>
		278	01000 10110	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SPRG7	SPR general 7	263	01000 00111	Read only	No	<a href="#">Section 3.10 on page 92</a>
		279	01000 10111	Read/Write	Yes	<a href="#">Section 3.10 on page 92</a>
SRR0	Save/restore register 0	26	00000 11010	Read/Write	Yes	<a href="#">Section 3.9.1.1 on page 83</a>
SRR1	Save/restore register 1	27	00000 11011	Read/Write	Yes	<a href="#">Section 3.9.1.2 on page 84</a>

**Table 96. Book E special-purpose registers (by SPR abbreviation) (continued)**

SPR	Name	Defined SPR number		Access	Supervisor only	Section/page
		Decimal	Binary			
TBL	Time base lower	268	01000 01100	Read only	No	<a href="#">Section 3.8.3 on page 81</a>
		284	01000 11100	Write-only	Yes	<a href="#">Section 3.8.3 on page 81</a>
TBU	Time base upper	269	01000 01101	Read only	No	<a href="#">Section 3.8.3 on page 81</a>
		285	01000 11101	Write-only	Yes	<a href="#">Section 3.8.3 on page 81</a>
TCR	Timer control register	340	01010 10100	Read/Write	Yes	<a href="#">Section 3.8.1 on page 79</a>
TSR	Timer status register	336	01010 10000	Read/Clear <sup>(4)</sup>	Yes	<a href="#">Section 3.8.2 on page 80</a>
USPRG0	User SPR general 0 <sup>(5)</sup>	256	01000 00000	Read/Write	No	<a href="#">Section 3.10 on page 92</a>
XER	Integer exception register	1	00000 00001	Read/Write	No	<a href="#">Section 3.3.2 on page 57</a>

1. The DBSR is read using **mf spr**. It cannot be directly written to. Instead, DBSR bits corresponding to 1 bits in the GPR can be cleared using **mt spr**.
2. Implementations may support more than one PID. If multiple PIDs are implemented, the Book E–defined PID is implemented as PID0.
3. User-mode read access to SPRG3 is implementation-dependent.
4. The TSR is read using **mf spr**. It cannot be directly written to. Instead, TSR bits corresponding to 1 bits in the GPR can be cleared using **mt spr**.
5. USPRG0 is a separate physical register from SPRG0.

[Table 97](#) lists EIS-specific SPRs, indicating which can be accessed by user-level software. Compilers should recognize SPR names when parsing instructions.

**Table 97. Implementation-specific SPRs (by SPR abbreviation)**

SPR	Name	SPR number	Access	Supervisor only	Section/page
ATBL	Alternate time base lower	526	Read-only	No	<a href="#">Section 3.15 on page 123</a>
ATBU	Alternate time base upper	527	Read-only	No	<a href="#">Section 3.15 on page 123</a>
DSRR0	Debug save/restore register 0	574	R/W	Yes	<a href="#">Section 3.9.1.10 on page 89</a>
DSRR1	Debug save/restore register 1	575	R/W	Yes	<a href="#">Section 3.9.1.10 on page 89</a>
IVOR32	SPE/embedded floating-point APU unavailable interrupt offset	528	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>





**Table 97. Implementation-specific SPRs (by SPR abbreviation) (continued)**

SPR	Name	SPR number	Access	Supervisor only	Section/page
IVOR33	Embedded floating-point data exception interrupt offset	529	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR34	Embedded floating-point round exception interrupt offset	530	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
IVOR35	Performance monitor	531	Read/Write	Yes	<a href="#">Section 3.9.1.7 on page 85</a>
L1CFG0	L1 cache configuration register 0	515	Read-only	No	<a href="#">Section 3.11.3 on page 96</a>
L1CFG1	L1 cache configuration register 1	516	Read-only	No	<a href="#">Section 3.11.3 on page 96</a>
L1CSR0	L1 cache control and status register 0	1010	Read/Write	Yes	<a href="#">Section 3.11.1 on page 92</a>
L1CSR1	L1 cache control and status register 1	1011	Read/Write	Yes	<a href="#">Section 3.11.2 on page 94</a>
L1FINV0	L1 flush and invalidate control register 0	1016	Read/Write	Yes	<a href="#">Section 3.11.5 on page 98</a>
MAS0	MMU assist register 0	624	Read/Write	Yes	<a href="#">Section 3.12.5.1 on page 103</a>
MAS1	MMU assist register 1	625	Read/Write	Yes	<a href="#">Section 3.12.5.1 on page 103</a>
MAS2	MMU assist register 2	626	Read/Write	Yes	<a href="#">Section 3.12.5.1 on page 103</a>
MAS3	MMU assist register 3	627	Read/Write	Yes	<a href="#">Section 3.12.5.4 on page 106</a>
MAS4	MMU assist register 4	628	Read/Write	Yes	<a href="#">Section 3.12.5.4 on page 106</a>
MAS5	MMU assist register 5.	629	Read/Write	Yes	<a href="#">Section 3.12.5.4 on page 106</a>
MAS6	MMU assist register 6	630	Read/Write	Yes	<a href="#">Section 3.12.5.4 on page 106</a>
MAS7	MMU assist register 7	944	Read/Write	Yes	<a href="#">Section 3.12.5.8 on page 109</a>
MCAR	Machine check address register	573	Read-only	Yes	<a href="#">Section 3.12.5.8 on page 109</a>
MCSR	Machine check syndrome register	572	Read/Write	Yes	<a href="#">Section 3.9.1.15 on page 90</a>
MCSRR0	Machine-check save/restore register 0	570	Read/Write	Yes	<a href="#">Section 3.9.1.15 on page 90</a>
MCSRR1	Machine-check save/restore register 1	571	Read/Write	Yes	<a href="#">Section 3.9.1.15 on page 90</a>

**Table 97. Implementation-specific SPRs (by SPR abbreviation) (continued)**

SPR	Name	SPR number	Access	Supervisor only	Section/page
MMUCFG	MMU configuration register	1015	Read-only	Yes	<a href="#">Section 3.12.3 on page 101</a>
MMUCSR0	MMU control and status register 0	1012	Read/Write	Yes	<a href="#">Section 3.12.2 on page 100</a>
PID0	Process ID register 0. Book E defines only this PID register and refers to as PID, not PID0.	48	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
PID1	Process ID register 1	633	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
PID2	Process ID register 2	634	Read/Write	Yes	<a href="#">Section 3.12.1 on page 99</a>
SPEFSCR	Signal processing and embedded floating-point status and control register	512	Read/Write	No	<a href="#">Section 3.14.1 on page 120</a>
SVR	System version register	1023	Read-only	Yes	<a href="#">Section 3.7.5 on page 77</a>
TLB0CFG	TLB configuration register 0	688	Read-only	Yes	<a href="#">Section 3.12.4 on page 102</a>
TLB1CFG	TLB configuration register 1	689	Read-only	Yes	<a href="#">Section 3.12.4 on page 102</a>

**4.3.1.15 Memory synchronization instructions**

Memory synchronization instructions control the order in which memory operations complete with respect to asynchronous events and the order in which memory operations are seen by other mechanisms that access memory. See [Table 98](#) for a summary.

**Table 98. Memory synchronization instructions**

Name	Mnemonic	Syntax	EIS notes
Instruction synchronize	<b>isync</b>	—	<p>Refetch serializing. An <b>isync</b> waits for previous instructions (including any interrupts they generate) to complete before <b>isync</b> executes, which purges all instructions from the processor and refetches the next instruction. <b>isync</b> does not wait for pending stores in the store queue to complete. Any subsequent instruction sees all effects of instructions before the <b>isync</b>.</p> <p>Because it prevents execution of subsequent instructions until preceding instructions complete, if an <b>isync</b> follows a conditional branch that depends on the value returned by a preceding load, the load on which the branch depends is performed before any loads caused by instructions after the <b>isync</b> even if the effects of the dependency are independent of the value loaded (for example, the value is compared to itself and the branch tests selected, CR<math>n</math>[EQ]), and even if the branch target is the next sequential instruction to be executed.</p>

Table 98. Memory synchronization instructions (continued)

Name	Mnemonic	Syntax	EIS notes
Load word and reserve indexed	<b>lwarx</b>	rD,rA,rB	<b>lwarx</b> with <b>stwcx.</b> can emulate semaphore operations such as test and set, compare and swap, exchange memory, and fetch and add. Both instructions must use the same EA. Reservation granularity is implementation-dependent. Executing <b>lwarx</b> and <b>stwcx.</b> to a page marked write-through (WIMG = 10xx) or when the data cache is locked may cause a data storage interrupt. If the location is not word-aligned, an alignment interrupt occurs.
Memory barrier	<b>mbar</b>	MO	<b>mbar</b> provides a pipelined memory barrier. (Note that <b>mbar</b> uses the same opcode as <b>eieio</b> , which is not defined by Book E.) The behavior of <b>mbar</b> is affected by the MO field (bits 6–10) of the instruction. MO = 0— <b>mbar</b> behaves identically to <b>msync</b> . MO = 1— <b>mbar</b> is a weaker, faster memory barrier; see the user's manual for implementation-specific behavior.
Memory synchronize	<b>msync</b>	—	Provides an ordering function for the effects of all instructions executed by the processor executing the <b>msync</b> . Executing an <b>msync</b> ensures that all previous instructions complete before it completes and that no subsequent instructions are initiated until after it completes. It also creates a memory barrier, which orders the storage accesses associated with these instructions. <b>msync</b> cannot complete before storage accesses associated with previous instructions are performed. <b>msync</b> is execution synchronizing. Note the following: <b>msync</b> is used to ensure that all stores into a data structure caused by store instructions executed in a critical section of a program are performed with respect to another processor before the store that releases the lock is performed with respect to that processor. <b>mbar</b> is preferable in many cases. On ST Book E devices: Unlike a context-synchronizing operations, <b>msync</b> does not discard prefetched instructions.
Store word conditional indexed	<b>stwcx.</b>	rS,rA,rB	<b>lwarx</b> with <b>stwcx.</b> can emulate semaphore operations such as test and set, compare and swap, exchange memory, and fetch and add. Both instructions must use the same EA. Reservation granularity is implementation-dependent. Executing <b>lwarx</b> and <b>stwcx.</b> to a page marked write-through (WIMG = 10xx) or cache-inhibited (WIMG = 01xx) when the data cache is locked may cause a data storage interrupt. If the location is not word-aligned, an alignment interrupt occurs.

4.3.1.16 Atomic update primitives using **lwarx** and **stwcx.**

The **lwarx** and **stwcx.** instructions together permit atomic update of a memory location. Book E provides word and double word forms of each of these instructions. Described here is the operation of **lwarx** and **stwcx.**

A specified memory location that may be modified by other processors or mechanisms requires memory coherence. If the location is in write-through required or caching inhibited

memory, the implementation determines whether these instructions function correctly or cause the system data storage error handler to be invoked.

Note the following:

- The memory coherence required attribute on other processors and mechanisms ensures that their stores to the specified location will cause the reservation created by the **lwarx** to be cancelled.
- **Warning:** Support for load and reserve and store conditional instructions for which the specified location is in caching-inhibited memory is being phased out of Book E. It is likely not to be provided on future implementations. New programs should not use these instructions to access caching inhibited memory.

A **lwarx** instruction is a load from a word-aligned location with the following side effects.

- A reservation for a subsequent **stwcx.** instruction is created.
- The memory coherence mechanism is notified that a reservation exists for the location accessed by the **lwarx**.

The **stwcx.** is a store to a word-aligned location that is conditioned on the existence of the reservation created by the **lwarx** and on whether both instructions specify the same location. To emulate an atomic operation, both **lwarx** and **stwcx.** must access the same location. **lwarx** and **stwcx.** are ordered by a dependence on the reservation, and the program is not required to insert other instructions to maintain the order of memory accesses caused by these two instructions.

A **stwcx.** performs a store to the target location only if the location accessed by the **lwarx** that established the reservation has not been stored into by another processor or mechanism between supplying a value for the **lwarx** and storing the value supplied by the **stwcx.**. If the instructions specify different locations, the store is not necessarily performed. CR0 is modified to indicate whether the store was performed, as follows:

$CR0[LT,GT,EQ,SO] = 0b00 \parallel \text{store\_performed} \parallel XER[SO]$

If a **stwcx.** completes but does not perform the store because a reservation no longer exists, CR0 is modified to indicate that the **stwcx.** completed without altering memory.

A **stwcx.** that performs its store is said to succeed.

Examples using **lwarx** and **stwcx.** are given in [Appendix C: Programming examples](#).

A successful **stwcx.** to a given location may complete before its store has been performed with respect to other processors and mechanisms. As a result, a subsequent load or **lwarx** from the given location on another processor may return a stale value. However, a subsequent **lwarx** from the given location on the other processor followed by a successful **stwcx.** on that processor is guaranteed to have returned the value stored by the first processor's **stwcx.** (in the absence of other stores to the given location).

#### 4.3.1.16.1 Reservations

The ability to emulate an atomic operation using **lwarx** and **stwcx.** is based on the conditional behavior of **stwcx.**, the reservation set by **lwarx**, and the clearing of that reservation if the target location is modified by another processor or mechanism before the **stwcx.** performs its store.

A reservation is held on an aligned unit of real memory called a reservation granule. The size of the reservation granule is implementation-dependent, but is a multiple of 4 bytes for **lwarx**. The reservation granule associated with EA contains the real address to which the EA maps. ('real\_addr(EA)' in the RTL for the load and reserve and store conditional

instructions stands for ‘real address to which EA maps.’) When one processor holds a reservation and another processor performs a store, the first processor’s reservation is cleared if the store affects any bytes in the reservation granule.

*Note: One use of **lwarx** and **stwcx**. is to emulate a compare and swap primitive like that provided by the IBM System/370 compare and swap instruction, which checks only that the old and current values of the word being tested are equal, with the result that programs that use such a compare and swap to control a shared resource can err if the word has been modified and the old value is subsequently restored. The use of **lwarx** and **stwcx**. improves on such a compare and swap, because the reservation reliably binds **lwarx** and **stwcx**. together. The reservation is always lost if the word is modified by another processor or mechanism between the **lwarx** and **stwcx**., so the **stwcx**. never succeeds unless the word has not been stored into (by another processor or mechanism) since the **lwarx**.*

A processor has at most one reservation at any time. Book E states that a reservation is established by executing a **lwarx** and is lost (or may be lost, in the case of the fourth and fifth bullets) if any of the following occurs.

- The processor holding the reservation executes another **lwarx**; this clears the first reservation and establishes a new one.
- The processor holding the reservation executes any **stwcx**., regardless of whether the specified address matches that of the **lwarx**.
- Another processor executes a store or **dcbz** to the same reservation granule.
- Another processor executes a **dcbtst**, **dcbst**, or **dcbf** to the same reservation granule; whether the reservation is lost is undefined.
- Another processor executes a **dcba** to the reservation granule. The reservation is lost if the instruction causes the target block to be newly established in the data cache or to be modified; otherwise, whether the reservation is lost is undefined.
- Some other mechanism modifies a location in the same reservation granule.
- Other implementation-specific conditions may also cause the reservation to be cleared, See the core reference manual.

Interrupts are not guaranteed to clear reservations. (However, system software invoked by interrupts may clear reservations.)

In general, programming conventions must ensure that **lwarx** and **stwcx**. specify addresses that match; a **stwcx**. should be paired with a specific **lwarx** to the same location. Situations in which a **stwcx**. may erroneously be issued after some **lwarx** other than that with which it is intended to be paired must be scrupulously avoided. For example, there must not be a context switch in which the processor holds a reservation in behalf of the old context, and the new context resumes after a **lwarx** and before the paired **stwcx**.. The **stwcx**. in the new context might succeed, which is not what was intended by the programmer.

Such a situation must be prevented by issuing a **stwcx**. to a dummy writable word-aligned location as part of the context switch, thereby clearing any reservation established by the old context. Executing **stwcx**. to a word-aligned location is enough to clear the reservation, regardless of whether it was set by **lwarx**.

#### 4.3.1.16.2 Forward progress

Forward progress in loops that use **lwarx** and **stwcx**. is achieved by a cooperative effort among hardware, operating system software, and application software.

Book E guarantees one of the following when a processor executes a **lwarx** to obtain a reservation for location X and then a **stwcx**. to store a value to location X:

1. The **stwcx.** succeeds and the value is written to location X.
2. The **stwcx.** fails because some other processor or mechanism modified location X.
3. The **stwcx.** fails because the processor's reservation was lost for some other reason.

In cases <sup>(1)</sup> and <sup>(2)</sup>, the system as a whole makes progress in the sense that some processor successfully modifies location X. Case <sup>(3)</sup> covers reservation loss required for correct operation of the rest of the system. This includes cancellation caused by some other processor writing elsewhere in the reservation granule for X, as well as cancellation caused by the operating system in managing certain limited resources such as real memory or context switches. It may also include implementation-dependent causes of reservation loss.

An implementation may make a forward progress guarantee, defining the conditions under which the system as a whole makes progress. Such a guarantee must specify the possible causes of reservation loss in case <sup>(3)</sup>. Although Book E alone cannot provide such a guarantee, the conditions in cases <sup>(1)</sup> and <sup>(2)</sup> are necessary for a guarantee. An implementation and operating system can build on them to provide such a guarantee.

Note that Book E does not guarantee fairness. In competing for a reservation, two processors can indefinitely lock out a third.

#### 4.3.1.16.3 Reservation loss due to granularity

Lock words should be allocated such that contention for the locks and updates to nearby data structures do not cause excessive reservation losses due to false indications of sharing that can occur due to the reservation granularity.

A processor holding a reservation on any word in a reservation granule loses its reservation if some other processor stores anywhere in that granule. Such problems can be avoided only by ensuring that few such stores occur. This can most easily be accomplished by allocating an entire granule for a lock and wasting all but one word.

Reservation granularity may vary for each implementation. There are no architectural restrictions bounding the granularity implementations must support, so reasonably portable code must dynamically allocate aligned and padded memory for locks to guarantee absence of granularity-induced reservation loss.

#### 4.3.1.17 Memory control instructions

Memory control instructions can be classified as follows:

- User- and supervisor-level cache management instructions.
- Supervisor-level-only translation lookaside buffer management instructions

This section describes the user-level cache management instructions. See [Section 4.3.2.2: Supervisor-level memory control instructions](#), for information about supervisor-level cache and translation lookaside buffer management instructions.

This section does not describe the cache-locking APU instructions, which are described in [Section 4.6.4: Cache locking APU](#).

##### 4.3.1.17.1 Cache management instructions

Cache management instructions obey the sequential execution model except as described in the example in this section of managing coherence between the instruction and data caches.

In the instruction descriptions the statements. “this instruction is treated as a load” and “this instruction is treated as a store,” mean that the instruction is treated as a load from or a store to the addressed byte with respect to address translation, memory protection, and the memory access ordering done by **msync**, **mbar**, and the other means described in [Section 6.2.2.1: Memory access ordering.](#)”

If caches are combined, the same value should be given for an instruction cache attribute and the corresponding data cache attribute.

Each implementation provides an efficient way for software to ensure that all blocks that are considered to be modified in the data cache have been copied to main memory before the processor enters any power-saving mode in which data cache contents are not maintained. The means are described in the reference manual for the implementation.

It is permissible for an implementation to treat any or all of the cache touch instructions (**icbt**, **dcbt**, or **dcbtst**) as no-operations, even if a cache is implemented.

The instruction cache is not necessarily kept consistent with the data cache or with main memory. When instructions are modified, software must ensure that the instruction cache is made consistent with data memory and that the modifications are made visible to the instruction fetching mechanism. The following instruction sequence can be used to accomplish this when the instructions being modified are in memory that is memory coherence required and one program both modifies the instructions and executes them. (Additional synchronization is needed when one program modifies instructions that another program will execute.) In this sequence, location ‘instr’ is assumed to contain modified instructions.

```

dcbst          instr          # update block in main memory
msync
icbi           instr          # invalidate copy in instr cache
msync
isync          # discard prefetched instructions

```

*Note:* Because the optimal instruction sequence may vary between systems, many operating systems provide a system service to perform the function described above. As stated above, the EA is translated using translation resources used for data accesses, even though the block being invalidated was copied into the instruction cache based on translation resources used for instruction fetches.

#### 4.3.1.18 User-level cache instructions

The instructions listed in [Table 99](#) help user-level programs manage on-chip caches if they are implemented. The following sections describe how these operations are treated with respect to the caches. The EIS supports the following CT values, defined by the EIS:

- CT = 0 indicates the L1 cache.
- CT = 1 indicates the I/O cache. (Note that some versions of the e500 documentation refer to the I/O cache as a frontside L2 cache.)
- CT = 2 indicates a backside L2 cache.

As with other memory-related instructions, the effects of cache management instructions on memory are weakly-ordered. If the programmer must ensure that cache or other instructions have been performed with respect to all other processors and system mechanisms, an **msync** must be placed after those instructions.

[Section 4.6.4: Cache locking APU](#), describes cache-locking APU instructions.

**Table 99. User-level cache instructions**

Name	Mnemonic	Syntax	Descriptions
Data cache block allocate	<b>dcba</b>	rA,rB	<p>This instruction is treated as a store with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.</p> <p>A no-op occurs if the cache is disabled or locked, if the page is marked write-through or cache-inhibited, or if a TLB protection violation occurs.</p> <p>An implementation may chose to no-op the instruction.</p>
Data cache block flush	<b>dcbf</b>	rA,rB	<p>This instruction is treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.</p>
Data cache block set to zero	<b>dcbz</b>	rA,rB	<p>This instruction is treated as a store with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.</p> <p>If the block containing the byte addressed by EA is in the data cache, all bytes of the block are cleared. If the block containing the byte addressed by EA is not in the data cache and is in storage that is not caching inhibited, the block is established in the data cache without fetching the block from main storage and all bytes of the block are cleared.</p> <p>If the block containing the byte addressed by EA is not in the data cache and is in storage that is not caching inhibited and cannot be established in the cache, then one of the following occurs:                      All bytes of the area of main storage that corresponds to the addressed block are set to zero                      An alignment interrupt is taken</p> <p>If the block containing the byte addressed by EA is in storage that is caching inhibited or write through required, one of the following occurs:                      All bytes of the area of main storage that corresponds to the addressed block are set to zero                      An alignment interrupt is taken.</p>
Data cache block store	<b>dcbst</b>	rA,rB	<p>This instruction is treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.</p>
Data cache block touch <sup>(1)</sup>	<b>dcbt</b>	CT,rA,rB	<p>This instruction is treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.</p> <p>A no-op occurs if the cache is disabled or locked, if the page is marked write-through or cache-inhibited, or if a TLB protection violation occurs.</p> <p>An implementation may chose to no-op the instruction.</p>



**Table 99. User-level cache instructions (continued)**

Name	Mnemonic	Syntax	Descriptions
Data cache block touch for store <sup>(1)</sup>	<b>dcbtst</b>	CT,rA,rB	Depending on the implementation, this instruction is treated as a load or store with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons. A no-op occurs if the cache is disabled or locked, if the page is marked write-through or cache-inhibited, or if a TLB protection violation occurs. An implementation may chose to no-op the instruction.
Instruction cache block invalidate	<b>icbi</b>	rA,rB	This instruction is treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.
Instruction cache block touch	<b>icbt</b>	CT,rA,rB	This instruction is treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons. A no-op occurs if the cache is disabled or locked, if the page is marked write-through or cache-inhibited, or if a TLB protection violation occurs. An implementation may chose to no-op the instruction.

1. A program that uses **dcbt** and **dcbtst** improperly is less efficient. To improve performance, **HID0[NOPTI]** can be set, which causes **dcbt** and **dcbtst** to be no-oped at the cache. They do not cause bus activity and cause only a 1-clock execution latency. The default state of this bit is zero, which enables the use of these instructions.

### 4.3.2 Supervisor level instructions

The Book E architecture includes the structure of the memory management model, supervisor-level registers, and the interrupt model. This section describes the supervisor-level instructions defined by the EIS.

#### 4.3.2.1 System linkage instructions

This section describes the system linkage instructions (see [Table 100](#)). The user-level **sc** instruction lets a user program call on the system to perform a service and causes the processor to take a system call interrupt. The supervisor-level **rfi** instruction is used for returning from an interrupt handler. The **rfdi** instruction is used for critical interrupts. The EIS defines the **rfmci** for machine check interrupts and **rfdi** for debug APU interrupts.

**Table 100. System linkage instructions—supervisor-level**

Name	Mnemonic	Syntax	Implementation notes
Return from interrupt	<b>rfi</b>	—	<b>rfi</b> is context-synchronizing
Return from debug interrupt	<b>rfdi</b>	—	Debug interrupt APU. When <b>rfdi</b> is executed, the values in the debug save and restore registers (DSRR0 and DSRR1) are restored. <b>rfdi</b> is context-synchronizing.
Return from machine check interrupt	<b>rfmci</b>	—	Machine check interrupt APU. When <b>rfmci</b> is executed, the values in the machine check interrupt save and restore registers (MCSRR0 and MCSRR1) are restored. <b>rfmci</b> is context-synchronizing.

**Table 100. System linkage instructions—supervisor-level (continued)**

Name	Mnemonic	Syntax	Implementation notes
Return from critical interrupt	<b>rfci</b>	—	When <b>rfci</b> executes, the values in the critical interrupt save and restore registers (CSRR0 and CSRR1) are restored. <b>rfci</b> is context-synchronizing.
System call	<b>sc</b>	—	The <b>sc</b> instruction is context-synchronizing.

[Table 101](#) lists instructions for accessing the MSR.

**Table 101. Move to/from machine state register instructions**

Name	Mnemonic	Syntax	Description
Move from machine state register	<b>mfmsr</b>	rD	—
Move to machine state register	<b>mtmsr</b>	rS	—
Write MSR external enable	<b>wrtee</b>	rS	Bit 48 of the contents of rS is placed into MSR[EE]. Other MSR bits are unaffected.
Write MSR external enable immediate	<b>wrteei</b>	E	The value of E is placed into MSR[EE]. Other MSR bits are unaffected.

Certain encodings of the SPR field of **mtspr** and **mfspir** instructions (shown in [Table 95](#)) provide access to supervisor-level SPRs. [Table 96](#) lists encodings for architecture-defined SPRs. Encodings for EIS-defined, supervisor-level SPRs are listed in [Table 102](#). Simplified mnemonics are provided for **mtspir** and **mfspir**. [Appendix C: Programming examples](#), describes context synchronization requirements when altering certain SPRs.

#### 4.3.2.2 Supervisor-level memory control instructions

Memory control instructions include the following:

- Cache management instructions (supervisor-level and user-level)
- Translation lookaside buffer management instructions

This section describes supervisor-level memory control instructions. [Section 4.3.1.17: Memory control instructions](#), describes user-level memory control instructions.

#### 4.3.2.3 Supervisor-level cache instruction

[Table 102](#) lists the only supervisor-level cache management instruction.

**Table 102. Supervisor-Level cache management instruction**

Name	Mnemonic	Syntax	Implementation notes
Data cache block invalidate	<b>dcbi</b>	rA,rB	This instruction is treated as a store with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons. An implementation may first perform a <b>dcbst</b> operation before invalidating the cache block if the memory is marked as coherency required (WIMG = xx1x).

See [Section 4.3.1.18: User-level cache instructions](#), for cache instructions that provide user-level programs the ability to manage the on-chip caches.

#### 4.3.2.4 Supervisor-level tlb management instructions

The address translation mechanism is defined in terms of TLBs and page table entries (PTEs) Book E processors use to locate the logical-to-physical address mapping for a particular access. See [Section 6.4: Storage model](#), for more information about TLB operations. [Table 103](#) summarizes the operation of the TLB instructions.

**Table 103. TLB management instructions**

Name	Mnemonic	Syntax	Implementation Notes
TLB invalidate virtual address indexed	tlbivax	rA, rB	A TLB invalidate operation is performed whenever <b>tlbivax</b> is executed. <b>tlbivax</b> invalidates any TLB entry that corresponds to the virtual address calculated by this instruction as long as IPROT is not set; this includes invalidating TLB entries contained in TLBs on other processors and devices in addition to the processor executing <b>tlbivax</b> . Thus an invalidate operation is broadcast throughout the coherent domain of the processor executing <b>tlbivax</b> . See <a href="#">Section 6.4</a> .
TLB read entry	tlbre	—	<b>tlbre</b> causes the contents of a single TLB entry to be extracted from the MMU and be placed in the corresponding MAS register fields. The entry extracted is specified by the TLBSEL, ESEL and EPN fields of MAS0 and MAS2. The contents extracted from the MMU are placed in MAS0–MAS3 and MAS7. See <a href="#">Section 6.4.9</a> .
TLB search indexed	tlbsx	rA, rB	<b>tlbsx</b> updates MAS conditionally based on the success or failure of a lookup in the MMU. The lookup is controlled by the EA provided by GPR[rB] specified in the instruction encoding and MAS6[SAS,SPID]. The values placed into MAS0–MAS3 and MAS7 differ, depending on whether a successful or unsuccessful search occurred. Note that RA=0 is a preferred form for <b>tlbsx</b> and that some ST implementations take an illegal instruction exception program interrupt if RA != 0.

**Table 103. TLB management instructions (continued)**

Name	Mnemonic	Syntax	Implementation Notes
TLB synchronize	tlbsync	—	<p>Provides an ordering function for the effects of all <b>tlbivax</b> instructions executed by the processor executing the <b>tlbsync</b> instruction, with respect to the memory barrier created by a subsequent <b>msync</b> instruction executed by the same processor. Executing a <b>tlbsync</b> instruction ensures that all of the following occurs:</p> <p>All TLB invalidations caused by <b>tlbivax</b> instructions preceding the <b>tlbsync</b> will have completed on any other processor before any storage accesses associated with data accesses caused by instructions following the <b>msync</b> instruction are performed with respect to that processor.</p> <p>All storage accesses by other processors for which the address was translated using the translations being invalidated, will have been performed with respect to the processor executing the <b>msync</b> instruction, to the extent required by the associated memory coherence required attributes, before the <b>mbar</b> or <b>msync</b> instruction's memory barrier is created.</p> <p>See <a href="#">Section 6.4.9: Translation lookaside buffer (TLB) arrays</a>.</p>
TLB Write Entry	tlbwe	—	<p><b>tlbwe</b> causes the contents of certain fields of MAS0, MAS1, MAS2, and MAS3 to be written into a TLB entry specified by the TLBSEL, ESEL, and EPN fields of MAS0 and MAS2. If MAS7 is implemented, execution of <b>tlbwe</b> causes any MAS7[RPN] to be written to the selected TLB entry. See <a href="#">Section 6.4.9</a>.</p>

### 4.3.3 Recommended simplified mnemonics

The description of each instruction includes the mnemonic and a formatted list of operands. Book E–compliant assemblers support the mnemonics and operand lists. To simplify assembly language programming, a set of simplified mnemonics and symbols is provided for some of the most frequently used instructions; refer to [Appendix B: Simplified mnemonics for PowerPC instructions](#), for a complete list. Programs written to be portable across the various assemblers for the Book E architecture should not assume the existence of mnemonics not described in this document.

### 4.3.4 Book E instructions with implementation-specific features

Book E defines several instructions in a general way, leaving the details of the execution up to the implementation. These are listed in [Table 104](#). This section describes how the EIS further defines those instructions. See the user documentation for additional implementation-specific behavior.

**Table 104. Implementation-specific instructions summary**

Name	Mnemonic	Syntax	Category
Move from APID Indirect	<b>mfapidi</b>	—	Optional. If not implemented, attempted execution causes an illegal instruction exception type program interrupt.
Move from Device Control Register	<b>mfdcr</b>	—	
Move to Device Control Register	<b>mtdcr</b>	—	

**Table 104. Implementation-specific instructions summary (continued)**

Name	Mnemonic	Syntax	Category
TLB Invalidate Virtual Address Indexed	<b>tlbivax</b>	rA, rB	These are described generally in <a href="#">Section 4.3.2.4: Supervisor-level tlb management instructions</a> .
TLB Read Entry	<b>tlbre</b>	—	
TLB Search Indexed	<b>tlbsx</b>	rA, rB	
TLB Write Entry	<b>tlbwe</b>	—	

A list of user-level instructions defined by both the classic PowerPC architecture and Book E can be found in [Section 4.7](#).

### 4.3.5 EIS instructions

The EIS defines the instructions listed in [Table 105](#) (with cross references to more detailed descriptions) that extend the Book E instruction set in accordance with Book E. SPE and embedded floating-point APU instructions are listed in [Table 108](#) and [Table 117](#).

**Table 105. EIS-defined instructions (except SPE and SPFP instructions)**

Name	Mnemonic	Syntax	Section #/page
Data Cache Block Lock Clear	<b>dcblc</b>	CT, rA, rB	<a href="#">Section 4.6.4</a>
Data Cache Block Touch and Lock Set	<b>dcbtls</b>	CT, rA, rB	
Data Cache Block Touch for Store and Lock Set	<b>dcbtstls</b>	CT, rA, rB	
Instruction Cache Block Lock Clear	<b>icblc</b>	CT, rA, rB	
Instruction Cache Block Touch and Lock Set	<b>icbtls</b>	CT, rA, rB	
Integer Select	<b>isel</b>	rD, rA, rB, crb	<a href="#">Section 4.6.2</a>
Move from Performance Monitor Register	<b>mfpmr</b>	rD, PMRN	<a href="#">Section 4.6.3</a>
Move to Performance Monitor Register	<b>mtpmr</b>	PMRN, rS	
Return from Machine Check Interrupt	<b>rfmci</b>	—	<a href="#">Section 4.6.5</a>
Return from Debug Interrupt	<b>rfdi</b>	—	<a href="#">Section 4.6.5</a>

### 4.3.6 Context synchronization

Context synchronization is achieved by post- and presynchronizing instructions. An instruction is presynchronized by completing all instructions before dispatching the presynchronized instruction. Post-synchronizing is implemented by not dispatching any later instructions until the post-synchronized instruction is completely finished.

## 4.4 Instruction fetching

In general, instructions are prefetched from the cache on a cache hit and from memory on a cache miss. Prefetched instructions may not be executed if the instruction stream is redirected after instructions are fetched and before they are scheduled for execution.

## 4.5 Memory synchronization

The **msync** instruction provides a memory barrier throughout the memory hierarchy. It waits for preceding data memory accesses to reach the point of coherency (that is, visible to the entire memory hierarchy); then it is broadcast. No subsequent instructions in the stream are initiated until after **msync** completes. Note that **msync** uses the same opcode as the **sync** instruction. The **msync** instruction is described in [Section 4.3.1.15: Memory synchronization instructions](#).

See [Section 6.2.2.1: Memory access ordering](#), for detailed information.

## 4.6 EIS-specific instructions

This section described EIS-defined instructions that are part of APUs or other extensions to the Book E architecture.

### 4.6.1 SPE and embedded floating-point APUs

The SPE and the embedded vector single-precision and embedded scalar double-precision APUs provide an extended GPR file with 32, 64-bit registers. The 32-bit Book E instructions operate on the lower (least significant) 32 bits of the 64-bit register. SPE APU vector instructions and embedded vector SPFP treat 64-bit registers as containing two 32-bit elements or four 16-bit elements as described in [Section 4.6.1.1.1: SPE APU instructions](#). The embedded double-precision floating-point APU uses the extended GPRs to hold single, IEEE-compliant double-precision operands.

However, like 32-bit Book E instructions, scalar SPFP APU floating-point instructions use bits 32–63 of the GPRs to hold 32-bit single-precision operands, as described in [Section 4.6.1.2: Embedded vector and scalar floating-point APU instructions](#).

There is no record form of SPE or embedded floating-point instructions. Vector compare instructions store the result of the comparison into the CR. The meaning of the CR bits is now overloaded for vector operations. Vector compare instructions specify a CR field and two source registers as well as the type of compare: greater than, less than, or equal. Two bits in the CR field are written with the result of the vector compare, one for each element. The two defined bits could be used either by a vector select instruction or by a UISA branch instruction.

A partially visible accumulator register is architected for the integer and fractional multiply accumulate SPE instructions. It is described in [Section 3.14.2](#).

Full descriptions of these instructions can be found in [Chapter 14: VLE instruction set](#).

#### 4.6.1.1 SPE APU instruction architecture

This section describes the instruction formats and instructions defined by the SPE APU.

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:

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

The real value of the binary operand SF[0:N-1] is as follows:

$$SF = -1.0 \cdot SF(0) + \sum_{i=1}^{N-1} SF(i) \cdot 2^{-i}$$

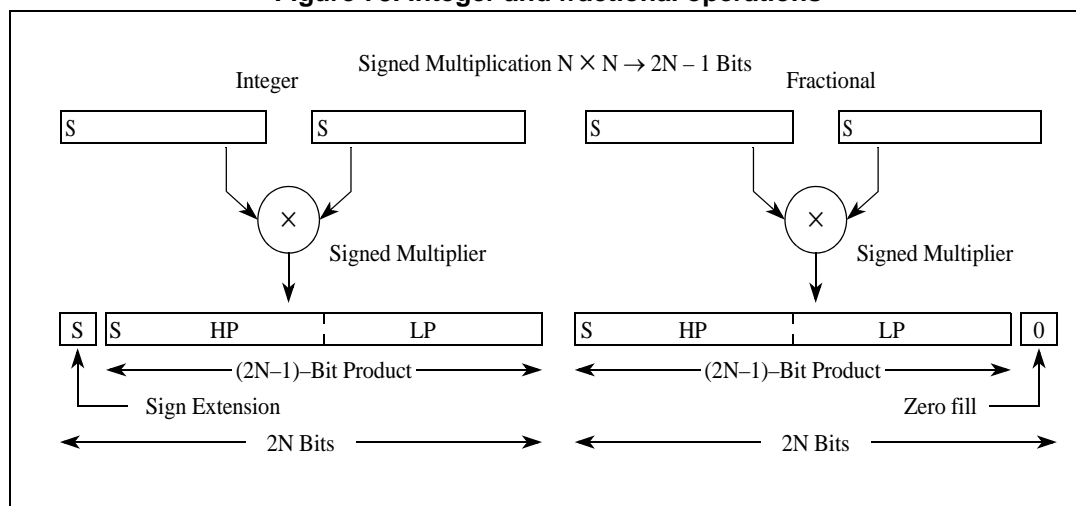
The most negative and positive numbers representable in fractional format are as follows:

- The most negative number is represented by SF(0) = 1 and SF[1:N-1] = 0 (that is, N=32; 0x8000\_0000 = -1.0).
- The most positive number is represented by SF(0) = 0 and SF[1:N-1] = all 1s (that is, N=32; 0x7FFF\_FFFF = 1.0 - 2<sup>-(N-1)</sup>).

SPE APU—integer and fractional operations

Figure 78 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 78. Integer and fractional operations



#### 4.6.1.1.1 SPE APU instructions

SPE APU instructions treat 64-bit GPRs as being composed of a vector of two 32-bit elements. (Some instructions also read or write 16-bit elements.) The SPE APU supports a number of forms of multiply and multiply-accumulate operations, and of add and subtract to accumulator operations. The SPE supports signed and unsigned forms, and optional fractional forms. For these instructions, the fractional form does not apply to unsigned forms because integer and fractional forms are identical for unsigned operands.

Table 106 shows how SPE APU vector multiply instruction mnemonics are structured.

**Table 106. SPE APU vector multiply instruction mnemonic structure**

Prefix	Multiply element		Data Type element		Accumulate element	
evm	ho	half odd (16x16->32)	usi umi ssi <b>ssf</b> <sup>(1)</sup> smi <b>smf</b> <sup>(1)</sup>	unsigned saturate integer unsigned modulo integer signed saturate integer signed saturate fractional signed modulo integer signed modulo fractional	a aa an aa w an w	write to ACC write to ACC & added ACC write to ACC & negate ACC write to ACC & ACC in words write to ACC & negate ACC in words
	he	half even (16x16->32)				
	hog	half odd guarded (16x16->32)				
	heg	half even guarded (16x16->32)				
	wh	word high (32x32->32)				
	wl	word low (32x32->32)				
	whg	word high guarded (32x32->32)				
	wlg	word low guarded (32x32->32)				
	w	word (32x32->64)				

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 107](#) defines mnemonic extensions for these instructions.

**Table 107. Mnemonic extensions for multiply-accumulate instructions**

Extension	Meaning	Comments
<b>Multiply form</b>		
<b>he</b>	Half word even	16x16→32
<b>heg</b>	Half word even guarded	16x16→32, 64-bit final accumulator result
<b>ho</b>	Half word odd	16x16→32
<b>hog</b>	Half word odd guarded	16x16→32, 64-bit final accumulator result
<b>w</b>	Word	32x32→64
<b>wh</b>	Word high	32x32→32, high-order 32 bits of product
<b>wl</b>	Word low	32x32→32, low-order 32 bits of product
<b>Data type</b>		
<b>smf</b>	Signed modulo fractional	(Wrap, no saturate)
<b>smi</b>	Signed modulo integer	(Wrap, no saturate)
<b>ssf</b>	Signed saturate fractional	
<b>ssi</b>	Signed saturate integer	
<b>umi</b>	Unsigned modulo integer	(Wrap, no saturate)
<b>usi</b>	Unsigned saturate integer	
<b>Accumulate options</b>		
<b>a</b>	Update accumulator	Update accumulator (no add)
<b>aa</b>	Add to accumulator	Add result to accumulator (64-bit sum)



**Table 107. Mnemonic extensions for multiply-accumulate instructions**

Extension	Meaning	Comments
<b>aaw</b>	Add to accumulator (words)	Add word results to accumulator words (pair of 32-bit sums)
<b>an</b>	Add negated	Add negated result to accumulator (64-bit sum)
<b>anw</b>	Add negated to accumulator (words)	Add negated word results to accumulator words (pair of 32-bit sums)

[Table 108](#) lists SPE APU instructions.

**Table 108. SPE APU vector instructions**

Instruction	Mnemonic	Syntax
Bit Reversed Increment	<b>brinc</b>	rD,rA,rB
Initialize Accumulator	<b>evmra</b>	rD,rA
Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate	<b>evmhsgmfaa</b>	rD,rA,rB
Multiply Half Words, Even, Guarded, Signed, Modulo, Fractional and Accumulate Negative	<b>evmhsgmfan</b>	rD,rA,rB
Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate	<b>evmhsgmiaa</b>	rD,rA,rB
Multiply Half Words, Even, Guarded, Signed, Modulo, Integer and Accumulate Negative	<b>evmhsgmian</b>	rD,rA,rB
Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate	<b>evmhgumiaa</b>	rD,rA,rB
Multiply Half Words, Even, Guarded, Unsigned, Modulo, Integer and Accumulate Negative	<b>evmhgumian</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate	<b>evmhogsmfaa</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Signed, Modulo, Fractional and Accumulate Negative	<b>evmhogsmfan</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer and Accumulate	<b>evmhogsmiaa</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Signed, Modulo, Integer and Accumulate Negative	<b>evmhogsmian</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate	<b>evmhogumiaa</b>	rD,rA,rB
Multiply Half Words, Odd, Guarded, Unsigned, Modulo, Integer and Accumulate Negative	<b>evmhogumian</b>	rD,rA,rB
Vector Absolute Value	<b>evabs</b>	rD,rA
Vector Add Immediate Word	<b>evaddiw</b>	rD,rB,UIMM
Vector Add Signed, Modulo, Integer to Accumulator Word	<b>evaddsmiaaw</b>	rD,rA,rB
Vector Add Signed, Saturate, Integer to Accumulator Word	<b>evaddssiaaw</b>	rD,rA
Vector Add Unsigned, Modulo, Integer to Accumulator Word	<b>evaddumiaaw</b>	rD,rA
Vector Add Unsigned, Saturate, Integer to Accumulator Word	<b>evaddusiaaw</b>	rD,rA

Table 108. SPE APU vector instructions (continued)

Instruction	Mnemonic	Syntax
Vector Add Word	<b>evaddw</b>	rD,rA,rB
Vector AND	<b>evand</b>	rD,rA,rB
Vector AND with Complement	<b>evandc</b>	rD,rA,rB
Vector Compare Equal	<b>evcmpeq</b>	crD,rA,rB
Vector Compare Greater Than Signed	<b>evcmpgts</b>	crD,rA,rB
Vector Compare Greater Than Unsigned	<b>evcmpgtu</b>	crD,rA,rB
Vector Compare Less Than Signed	<b>evcmplt</b>	crD,rA,rB
Vector Compare Less Than Unsigned	<b>evcmpltu</b>	crD,rA,rB
Vector Convert Floating-Point from Signed Fraction	<b>evfscfsf</b>	rD,rB
Vector Convert Floating-Point from Signed Integer	<b>evfscfsi</b>	rD,rB
Vector Convert Floating-Point from Unsigned Fraction	<b>evfscfuf</b>	rD,rB
Vector Convert Floating-Point from Unsigned Integer	<b>evfscfui</b>	rD,rB
Vector Convert Floating-Point to Signed Fraction	<b>evfsctsf</b>	rD,rB
Vector Convert Floating-Point to Signed Integer	<b>evfsctsi</b>	rD,rB
Vector Convert Floating-Point to Signed Integer with Round toward Zero	<b>evfsctsiz</b>	rD,rB
Vector Convert Floating-Point to Unsigned Fraction	<b>evfsctuf</b>	rD,rB
Vector Convert Floating-Point to Unsigned Integer	<b>evfsctui</b>	rD,rB
Vector Convert Floating-Point to Unsigned Integer with Round toward Zero	<b>evfsctuiz</b>	rD,rB
Vector Count Leading Sign Bits Word	<b>evcntlsw</b>	rD,rA
Vector Count Leading Zeros Word	<b>evcntlzw</b>	rD,rA
Vector Divide Word Signed	<b>evdivws</b>	rD,rA,rB
Vector Divide Word Unsigned	<b>evdivwu</b>	rD,rA,rB
Vector Equivalent	<b>eveqv</b>	rD,rA,rB
Vector Extend Sign Byte	<b>evextsb</b>	rD,rA
Vector Extend Sign Half Word	<b>evextsh</b>	rD,rA
Vector Floating-Point Absolute Value	<b>evfsabs</b>	rD,rA
Vector Floating-Point Add	<b>evfsadd</b>	rD,rA,rB
Vector Floating-Point Compare Equal	<b>evfscmpeq</b>	crD,rA,rB
Vector Floating-Point Compare Greater Than	<b>evfscmpgt</b>	crD,rA,rB
Vector Floating-Point Compare Less Than	<b>evfscmplt</b>	crD,rA,rB
Vector Floating-Point Divide	<b>evfsdiv</b>	rD,rA,rB
Vector Floating-Point Multiply	<b>evfsmul</b>	rD,rA,rB
Vector Floating-Point Negate	<b>evfsneg</b>	rD,rA
Vector Floating-Point Negative Absolute Value	<b>evfsnabs</b>	rD,rA
Vector Floating-Point Subtract	<b>evfssub</b>	rD,rA,rB

Table 108. SPE APU vector instructions (continued)

Instruction	Mnemonic	Syntax
Vector Floating-Point Test Equal	evfststeq	crD,rA,rB
Vector Floating-Point Test Greater Than	evfststgt	crD,rA,rB
Vector Floating-Point Test Less Than	evfststlt	crD,rA,rB
Vector Load Double into Half Words	evldh	rD,d(rA)
Vector Load Double into Half Words Indexed	evldhx	rD,rA,rB
Vector Load Double into Two Words	evldw	rD,d(rA)
Vector Load Double into Two Words Indexed	evldwx	rD,rA,rB
Vector Load Double Word into Double Word	evldd	rD,d(rA)
Vector Load Double Word into Double Word Indexed	evlddx	rD,rA,rB
Vector Load Half Word into Half Word Odd Signed and Splat	evlhhosspat	rD,d(rA)
Vector Load Half Word into Half Word Odd Signed and Splat Indexed	evlhhosspatx	rD,rA,rB
Vector Load Half Word into Half Word Odd Unsigned and Splat	evlhhuspat	rD,d(rA)
Vector Load Half Word into Half Word Odd Unsigned and Splat Indexed	evlhhuspatx	rD,rA,rB
Vector Load Half Word into Half Words Even and Splat	evlhhespat	rD,d(rA)
Vector Load Half Word into Half Words Even and Splat Indexed	evlhhespatx	rD,rA,rB
Vector Load Word into Half Words and Splat	evlwspat	rD,d(rA)
Vector Load Word into Half Words and Splat Indexed	evlwspatx	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	evlwwspat	rD,d(rA)
Vector Load Word into Word and Splat Indexed	evlwwspatx	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
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

**Table 108. SPE APU vector instructions (continued)**

Instruction	Mnemonic	Syntax
Vector Multiply Half Words, Even, Signed, Modulo, Integer	<b>evmhesmi</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate into Words	<b>evmhesmiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Modulo, Integer and Accumulate Negative into Words	<b>evmhesmianw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Modulo, Integer, Accumulate	<b>evmhesmia</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Fractional	<b>evmhessf</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate into Words	<b>evmhessfaaw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Fractional and Accumulate Negative into Words	<b>evmhessfanw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Fractional, Accumulate	<b>evmhessfa</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate into Words	<b>evmhessiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Signed, Saturate, Integer and Accumulate Negative into Words	<b>evmhessianw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Modulo, Integer	<b>evmheumi</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate into Words	<b>evmheumiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Modulo, Integer and Accumulate Negative into Words	<b>evmheumianw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Modulo, Integer, Accumulate	<b>evmheumia</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate into Words	<b>evmheusiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Even, Unsigned, Saturate, Integer and Accumulate Negative into Words	<b>evmheusianw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Fractional	<b>evmhosmf</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate into Words	<b>evmhosmfaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Fractional and Accumulate Negative into Words	<b>evmhosmfanw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Fractional, Accumulate	<b>evmhosmfa</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Integer	<b>evmhosmi</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate into Words	<b>evmhosmiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Integer and Accumulate Negative into Words	<b>evmhosmianw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Modulo, Integer, Accumulate	<b>evmhosmia</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Saturate, Fractional	<b>evmhossf</b>	rD,rA,rB

**Table 108. SPE APU vector instructions (continued)**

Instruction	Mnemonic	Syntax
Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate into Words	<b>evmhossfaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Saturate, Fractional and Accumulate Negative into Words	<b>evmhossfanw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Saturate, Fractional, Accumulate	<b>evmhossfa</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate into Words	<b>evmhossiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Signed, Saturate, Integer and Accumulate Negative into Words	<b>evmhossianw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer	<b>evmhoumi</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate into Words	<b>evmhoumiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer and Accumulate Negative into Words	<b>evmhoumianw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Modulo, Integer, Accumulate	<b>evmhoumia</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate into Words	<b>evmhousiaaw</b>	rD,rA,rB
Vector Multiply Half Words, Odd, Unsigned, Saturate, Integer and Accumulate Negative into Words	<b>evmhousianw</b>	rD,rA,rB
Vector Multiply Word High Signed, Modulo, Fractional	<b>evmwhsmf</b>	rD,rA,rB
Vector Multiply Word High Signed, Modulo, Fractional and Accumulate	<b>evmwhsmfa</b>	rD,rA,rB
Vector Multiply Word High Signed, Modulo, Integer	<b>evmwhsmi</b>	rD,rA,rB
Vector Multiply Word High Signed, Modulo, Integer and Accumulate	<b>evmwhsmia</b>	rD,rA,rB
Vector Multiply Word High Signed, Saturate, Fractional	<b>evmwhssf</b>	rD,rA,rB
Vector Multiply Word High Signed, Saturate, Fractional and Accumulate	<b>evmwhssfa</b>	rD,rA,rB
Vector Multiply Word High Unsigned, Modulo, Integer	<b>evmwhumi</b>	rD,rA,rB
Vector Multiply Word High Unsigned, Modulo, Integer and Accumulate	<b>evmwhumia</b>	rD,rA,rB
Vector Multiply Word Low Signed, Modulo, Integer and Accumulate in Words	<b>evmwlsmiaaw</b>	rD,rA,rB
Vector Multiply Word Low Signed, Modulo, Integer and Accumulate Negative in Words	<b>evmwlsmianw</b>	rD,rA,rB
Vector Multiply Word Low Signed, Saturate, Integer and Accumulate in Words	<b>evmwlssiaaw</b>	rD,rA,rB
Vector Multiply Word Low Signed, Saturate, Integer and Accumulate Negative in Words	<b>evmwlssianw</b>	rD,rA,rB
Vector Multiply Word Low Unsigned, Modulo, Integer	<b>evmwlumi</b>	rD,rA,rB
Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate	<b>evmwlumia</b>	rD,rA,rB
Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate in Words	<b>evmwlumiaaw</b>	rD,rA,rB
Vector Multiply Word Low Unsigned, Modulo, Integer and Accumulate Negative in Words	<b>evmwlumianw</b>	rD,rA,rB

**Table 108. SPE APU vector instructions (continued)**

Instruction	Mnemonic	Syntax
Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate in Words	<b>evmwlusiaaw</b>	rD,rA,rB
Vector Multiply Word Low Unsigned, Saturate, Integer and Accumulate Negative in Words	<b>evmwlusianw</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Fractional	<b>evmwsmf</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Fractional and Accumulate	<b>evmwsmfa</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Fractional and Accumulate	<b>evmwsmfaa</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Fractional and Accumulate Negative	<b>evmwsmfan</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Integer	<b>evmwsmi</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Integer and Accumulate	<b>evmwsmia</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Integer and Accumulate	<b>evmwsmiaa</b>	rD,rA,rB
Vector Multiply Word Signed, Modulo, Integer and Accumulate Negative	<b>evmwsmian</b>	rD,rA,rB
Vector Multiply Word Signed, Saturate, Fractional	<b>evmwssf</b>	rD,rA,rB
Vector Multiply Word Signed, Saturate, Fractional and Accumulate	<b>evmwssfafa</b>	rD,rA,rB
Vector Multiply Word Signed, Saturate, Fractional and Accumulate	<b>evmwssfafa</b>	rD,rA,rB
Vector Multiply Word Signed, Saturate, Fractional and Accumulate Negative	<b>evmwssfafan</b>	rD,rA,rB
Vector Multiply Word Unsigned, Modulo, Integer	<b>evmwumi</b>	rD,rA,rB
Vector Multiply Word Unsigned, Modulo, Integer and Accumulate	<b>evmwumia</b>	rD,rA,rB
Vector Multiply Word Unsigned, Modulo, Integer and Accumulate	<b>evmwumiaa</b>	rD,rA,rB
Vector Multiply Word Unsigned, Modulo, Integer and Accumulate Negative	<b>evmwumian</b>	rD,rA,rB
Vector NAND	<b>evnand</b>	rD,rA,rB
Vector Negate	<b>evneg</b>	rD,rA
Vector NOR	<b>evnor</b>	rD,rA,rB
Vector OR	<b>evor</b>	rD,rA,rB
Vector OR with Complement	<b>evorc</b>	rD,rA,rB
Vector Rotate Left Word	<b>evrlw</b>	rD,rA,rB
Vector Rotate Left Word Immediate	<b>evrlwi</b>	rD,rA,UIMM
Vector Round Word	<b>evrndw</b>	rD,rA
Vector Select	<b>evsel</b>	rD,rA,rB,crS
Vector Shift Left Word	<b>evslw</b>	rD,rA,rB
Vector Shift Left Word Immediate	<b>evslwi</b>	rD,rA,UIMM
Vector Shift Right Word Immediate Signed	<b>evsrwis</b>	rD,rA,UIMM
Vector Shift Right Word Immediate Unsigned	<b>evsrwiu</b>	rD,rA,UIMM
Vector Shift Right Word Signed	<b>evsrws</b>	rD,rA,rB
Vector Shift Right Word Unsigned	<b>evsrwu</b>	rD,rA,rB

Table 108. SPE APU vector instructions (continued)

Instruction	Mnemonic	Syntax
Vector Splat Fractional Immediate	evsplatfi	rD,SIMM
Vector Splat Immediate	evsplat	rD,SIMM
Vector Store Double of Double	evstd	rS,d(rA)
Vector Store Double of Double Indexed	evstdx	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	evstwe	rS,d(rA)
Vector Store Word of Two Half Words from Even Indexed	evstweh	rS,rA,rB
Vector Store Word of Two Half Words from Odd	evstwo	rS,d(rA)
Vector Store Word of Two Half Words from Odd Indexed	evstwoh	rS,rA,rB
Vector Store Word of Word from Even	evstwe	rS,d(rA)
Vector Store Word of Word from Even Indexed	evstweh	rS,rA,rB
Vector Store Word of Word from Odd	evstwo	rS,d(rA)
Vector Store Word of Word from Odd Indexed	evstwoh	rS,rA,rB
Vector Subtract from Word	evsubfw	rD,rA,rB
Vector Subtract Immediate from Word	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
Vector Subtract Unsigned, Saturate, Integer to Accumulator Word	evsubfusiaaw	rD,rA
Vector XOR	evxor	rD,rA,rB

#### 4.6.1.2 Embedded vector and scalar floating-point APU instructions

The embedded floating-point operations are IEEE-compliant with software exception handlers and offer a simpler exception model than the floating-point instructions defined by the PowerPC ISA. 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.

The SPFP APUs are described as follows:

- Vector SPFP instructions operate on a vector of two 32-bit, single-precision floating-point numbers that reside in the upper and lower halves of the 64-bit GPRs. These instructions are listed in [Table 117](#) alongside their scalar equivalents.
- Scalar SPFP instructions operate on single 32-bit operands that reside in the lower 32-bits of the GPRs. These instructions are listed in [Table 117](#).
- Scalar DFP instructions operate on single 64-bit double-precision operands that reside in the 64-bit GPRs. These instructions are listed in [Table 109](#).

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

**Table 109. Vector and scalar floating-point APU instructions**

Instruction	Single-precision		Double-precision scalar	Syntax
	Scalar	Vector		
Convert Floating-Point Double- from Single-Precision	—	—	efdcfs	rD,rB
Convert Floating-Point from Signed Fraction	efscfsf	evscfsf	efdcfsf	rD,rB
Convert Floating-Point from Signed Integer	efscfsi	evscfsi	efdcfsi	rD,rB
Convert Floating-Point from Unsigned Fraction	efscfuf	evscfuf	efdcfuf	rD,rB
Convert Floating-Point from Unsigned Integer	efscfui	evscfui	efdcfui	rD,rB
Convert Floating-Point Single- from Double-Precision	—	—	efscfd	rD,rB
Convert Floating-Point to Signed Fraction	efscfsf	evscfsf	efdcfsf	rD,rB
Convert Floating-Point to Signed Integer	efscfsi	evscfsi	efdcfsi	rD,rB
Convert Floating-Point to Signed Integer with Round toward Zero	efscfsiz	evscfsiz	efdcfsiz	rD,rB
Convert Floating-Point to Unsigned Fraction	efscfuf	evscfuf	efdcfuf	rD,rB
Convert Floating-Point to Unsigned Integer	efscfui	evscfui	efdcfui	rD,rB
Convert Floating-Point to Unsigned Integer with Round toward Zero	efscfuiZ	evscfuiZ	efdcfuiZ	rD,rB
Floating-Point Absolute Value	efsabs <sup>(1)</sup>	evfsabs	efdabs	rD,rA
Floating-Point Add	efsadd	evfsadd	efdadd	rD,rA,rB
Floating-Point Compare Equal	efscmpeq	evscmpeq	efdcmpeq	crD,rA,rB
Floating-Point Compare Greater Than	efscmpgt	evscmpgt	efdcmpgt	crD,rA,rB
Floating-Point Compare Less Than	efscmplt	evscmplt	efdcmplt	crD,rA,rB
Floating-Point Divide	efdiv	evfdiv	efddiv	rD,rA,rB
Floating-Point Multiply	efsmul	evfsmul	efdmul	rD,rA,rB
Floating-Point Negate	efsneg <sup>(1)</sup>	evfsneg	efdneg	rD,rA
Floating-Point Negative Absolute Value	efsnabs <sup>(1)</sup>	evfsnabs	efdnabs	rD,rA
Floating-Point Subtract	efssub	evfssub	efdsb	rD,rA,rB
Floating-Point Test Equal	efststeq	evfststeq	efdststeq	crD,rA,rB
Floating-Point Test Greater Than	efststgt	evfststgt	efdststgt	crD,rA,rB



**Table 109. Vector and scalar floating-point APU instructions (continued)**

Instruction	Single-precision		Double-precision scalar	Syntax
	Scalar	Vector		
Floating-Point Test Less Than	efstslt	evfstslt	efdtslt	crD,rA,rB
<b>SPE Double Word Load/Store Instructions</b>				
Vector Load Double Word into Double Word	—	evldd	evldd	rD,d(rA)
Vector Load Double Word into Double Word Indexed	—	evlddx	evlddx	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

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.

### 4.6.2 Integer select (isel) APU

The integer select APU consists of the **isel** instruction, a conditional register move that helps eliminate branches. [Section 8.1: Integer select APU](#), describes the use of **isel**.

**Table 110. Integer select APU instruction**

Name	Mnemonic	Syntax
Integer Select	isel	rD,rA,rB,crB

### 4.6.3 Performance monitor APU

The EIS defines the performance monitor as an APU. Software communication with the performance monitor APU is achieved through performance monitor registers (PMRs) rather than SPRs. New instructions are provided to move to and from these PMRs. Performance monitor APU instructions are described in [Table 111](#).

**Table 111. Performance monitor APU instructions**

Name	Mnemonic	Syntax
Move from performance monitor register	mfpmr	rD,PMRN
Move to performance monitor register	mtpmr	PMRN,rS

The Book E implementation standards defines a set of register resources used exclusively by the performance monitor. PMRs are similar to the SPRs defined in the Book E architecture and are accessed by **mtpmr** and **mfpmr**, which are also defined by the EIS.

Table 112 lists supervisor-level PMRs. User-level software that attempts to read or write supervisor-level PMRs causes a privilege exception.

**Table 112. Performance monitor registers—supervisor level**

Abbreviation	Register name	PMR number	pmr[0–4]	pmr[5–9]	Section/page
PMGC0	Performance monitor global control register 0	400	01100	10000	<a href="#">Section 3.16.1 on page 126</a>
PMLCa0	Performance monitor local control a0	144	00100	10000	<a href="#">Section 3.16.3 on page 127</a>
PMLCa1	Performance monitor local control a1	145	00100	10001	
PMLCa2	Performance monitor local control a2	146	00100	10010	
PMLCa3	Performance monitor local control a3	147	00100	10011	
PMLCb0	Performance monitor local control b0	272	01000	10000	<a href="#">Section 3.16.5 on page 128</a>
PMLCb1	Performance monitor local control b1	273	01000	10001	
PMLCb2	Performance monitor local control b2	274	01000	10010	
PMLCb3	Performance monitor local control b3	275	01000	10011	
PMC0	Performance monitor counter 0	16	00000	10000	<a href="#">Section 3.16.7 on page 130</a>
PMC1	Performance monitor counter 1	17	00000	10001	
PMC2	Performance monitor counter 2	18	00000	10010	
PMC3	Performance monitor counter 3	19	00000	10011	

User-level PMRs in Table 113 are read-only and are accessed with **mfpmr**. Attempting to write user-level registers in supervisor or user mode causes an illegal instruction exception.

Table 113. Performance monitor registers—user level (read-only)

Abbreviation	Register Name	PMR Number	pmr[0–4]	pmr[5–9]	Section/Page
UPMGC0	User performance monitor global control register 0	384	01100	00000	<a href="#">Section 3.16.2 on page 127</a>
UPMLCa0	User performance monitor local control a0	128	00100	00000	<a href="#">Section 3.16.4 on page 128</a>
UPMLCa1	User performance monitor local control a1	129	00100	00001	
UPMLCa2	User performance monitor local control a2	130	00100	00010	
UPMLCa3	User performance monitor local control a3	131	00100	00011	
UPMLCb0	User performance monitor local control b0	256	01000	00000	<a href="#">Section 3.16.6 on page 130</a>
UPMLCb1	User performance monitor local control b1	257	01000	00001	
UPMLCb2	User performance monitor local control b2	258	01000	00010	
UPMLCb3	User performance monitor local control b3	259	01000	00011	
UPMC0	User performance monitor counter 0	0	00000	00000	<a href="#">Section 3.16.8 on page 130</a>
UPMC1	User performance monitor counter 1	1	00000	00001	
UPMC2	User performance monitor counter 2	2	00000	00010	
UPMC3	User performance monitor counter 3	3	00000	00011	

#### 4.6.4 Cache locking APU

This section describes the instructions in the cache locking APU, which consists of the instructions described in [Table 114](#). Lines are locked into the cache by software using a series of touch and lock set instructions. The following instructions are provided to lock data items into the data and instruction cache:

- **dcbtIs**—Data Cache Block Touch and Lock Set
- **dcbtstIs**—Data Cache Block Touch for Store and Lock Set
- **icbtIs**—Instruction Cache Block Touch and Lock Set

The **rA** and **rB** operands to these instructions form a EA identifying the line to be locked. The **CT** field indicates which cache in the cache hierarchy should be targeted. These instructions are similar to the **dcbt**, **dcbtst**, and **icbt** instructions, but locking instructions can not execute speculatively and may cause additional exceptions. For unified caches, both the instruction lock set and the data lock set target the same cache.

Similarly, lines are unlocked from the cache by software using a series of lock-clear instructions. The following instructions are provided to lock instructions into the instruction cache:

- **dcblc**—Data Cache Block Lock Clear
- **icblc**—Instruction Cache Block Lock Clear

The rA and rB operands to these instructions form an EA identifying the line to be unlocked. The CT field indicates which cache in the cache hierarchy should be targeted.

Additionally, software may clear all the locks in the cache. For the primary cache, this is accomplished by setting the CLFC (DCLFC, ICLFC) bit in L1CSR0 (L1CSR1).

Cache lines can also be implicitly unlocked in the following ways:

- A locked line is invalidated if it is targeted by a **dcbi**, **dcbf**, or **icbi** instruction.
- A snoop hit on a locked line that requires the line to be invalidated. This can occur because the data the line contains has been modified external to the processor, or another processor has explicitly invalidated the line.
- The entire cache containing the locked line is flash invalidated.

An implementation is not required to unlock lines if data is invalidated in the cache. Although the data may be invalidated (and thus not in the cache), the cache can keep the lock associated with that cache line present and fill the line from the memory subsystem when the next access occurs. If the implementation does not clear locks when the associated line is invalidated, the method of locking is said to be persistent. An implementation may choose to implement locks as persistent or not persistent; the preferred method is persistent.

**Table 114. Cache locking APU instructions**

Name	Mnemonic	Syntax	Description
Data Cache Block Lock Clear	<b>dcblc</b>	CT,rA,rB	Treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.
Data Cache Block Touch and Lock Set	<b>dcbtls</b>	CT,rA,rB	Treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.
Data Cache Block Touch for Store and Lock Set	<b>dcbtstls</b>	CT,rA,rB	It is implementation dependent whether this instruction is treated as a load or store with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.
Instruction Cache Block Lock Clear	<b>icblc</b>	CT,rA,rB	Treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.
Instruction Cache Block Touch and Lock Set	<b>icbtls</b>	CT,rA,rB	Treated as a load with respect to any memory barriers, synchronization, translation and protection, and debug address comparisons.

The cache-locking APU defines a flash clear for all data cache lock bits (using L1CSR0[CLFR]) and in the instruction cache (using L1CSR1[ICLFR]). This allows system software to clear all data cache locking bits without knowing the addresses of the lines locked.

## 4.6.5 Machine check APU

The machine check APU defines a separate interrupt type for machine check interrupts. It provides additional save and restore SPRs (MCSRR and MCSRR1). The Return from Machine Check Interrupt instruction (**rfmci**), is described in [Table 115](#).

**Table 115. Machine check APU instruction**

Name	Mnemonic	Syntax	Implementation notes
Return from machine check interrupt	<b>rfmci</b>	—	Restores MCSRR0 and MCSRR1 values; context-synchronizing.

## 4.6.6 VLE extension

This section lists instructions defined or supported by the VLE extension. Unless otherwise noted, instructions that are not prefixed with **e\_** or **se\_** have identical encodings and semantics as in Book E or in the EIS. Book E–defined instructions listed in the tables in this section can be executed when the processor is in VLE mode; Book E instructions not listed cannot.

A complete list of supported instructions is provided in [Section 4.6.6.11: Instruction listings](#).

### 4.6.6.1 Processor control instructions

This section lists processor control instructions that can be executed when a processor is in VLE mode. These instructions are grouped as follows:

- [Section 4.6.6.1.1: System linkage instructions](#)
- [Section 4.6.6.1.2: Processor control register manipulation instructions](#)
- [Section 4.6.6.2: Instruction synchronization instruction](#)

#### 4.6.6.1.1 System linkage instructions

The **se\_sc**, **se\_rfi**, **se\_rfci**, and **se\_rfdi** system linkage instructions, shown in [Table 116](#), enable a program to call on the system to perform a service (that is, invoke a system call interrupt), and enable the system to return from performing a service or from processing an interrupt.

**Table 116. System linkage instruction set index**

Mnemonic	Instruction	Reference
<b>se_sc</b>	System Call	<a href="#">Section on page 769</a>
<b>se_rfci`</b>	Return from critical interrupt	<a href="#">Section on page 767</a>
<b>se_rfdi</b>	Return from debug interrupt	<a href="#">Section on page 704</a>
<b>se_rfi</b>	Return from interrupt	<a href="#">Section on page 767</a>

#### 4.6.6.1.2 Processor control register manipulation instructions

In addition to the Book E processor control register manipulation instructions, the VLE extension provides 16-bit forms of instructions to move to/from the LR and CTR, listed in [Table 117](#)

**Table 117. System register manipulation instruction set index**

Mnemonic	Instruction	Reference
<b>se_mfctr</b> rX	Move From Count Register	<a href="#">Section on page 762</a>
<b>mfdcr</b> rD,DCRN	Move From Device Control Register	Book E
<b>se_mflr</b> rX	Move From Link Register	<a href="#">Section on page 762</a>
<b>mfmsr</b> rD	Move From Machine State Register	Book E
<b>mfspir</b> rD,SPRN	Move From Special Purpose Register	Book E
<b>se_mtctr</b> rX	Move To Count Register	<a href="#">Section on page 763</a>
<b>mtdcr</b> DCRN,rS	Move To Device Control Register	Book E
<b>se_mtlr</b> rX	Move To Link Register	<a href="#">Section on page 763</a>
<b>mtmsr</b> rS	Move To Machine State Register	Book E
<b>mtspr</b> SPRN,rS	Move To Special Purpose Register	Book E
<b>wrttee</b> rA	Write MSR External Enable	Book E
<b>wrtteei</b> E	Write MSR External Enable Immediate	Book E

#### 4.6.6.2 Instruction synchronization instruction

[Table 118](#) lists the VLE-defined **se\_isync** instruction.

**Table 118. Instruction Synchronization Instruction Set Index**

Mnemonic	Instruction	Reference
<b>se_isync</b>	Instruction Synchronize	<a href="#">Section on page 757</a>

#### 4.6.6.3 Branch operation instructions

This section lists branch instructions that can be executed when a processor is in VLE mode. It also describes the registers that support them.

Registers for branch operations

The sections listed in the following describe the registers that support branch operations:

- [Section 3.5.1: Condition register \(CR\)](#)
- [Section 3.5.2: Link register \(LR\)](#)
- [Section 3.5.3: Count register \(CTR\)](#)

#### 4.6.6.3.1 Branch instructions

The sequence of instruction execution can be changed by the branch instructions. Because VLE instructions must be aligned on half-word boundaries, the low-order bit of the generated branch target address is forced to 0 by the processor in performing the branch.

The branch instructions compute the EA of the target in one of the following ways, as described in [Section 11.2: Instruction memory addressing modes](#).

1. Adding a displacement to the address of the branch instruction.
2. Using the address contained in the LR (Branch to Link Register [and Link]).
3. Using the address contained in the CTR (Branch to Count Register [and Link]).

Branching can be conditional or unconditional, and the return address can optionally be provided. If the return address is to be provided (LK = 1), the EA of the instruction following the branch instruction is placed into the LR after the branch target address has been computed: this is done whether or not the branch is taken.

In branch conditional instructions, the BI32 or BI16 instruction field specifies the CR bit to be tested. For 32-bit instructions using BI32, CR[32–47] (corresponding to bits in CR0–CR3) may be specified. For 16-bit instructions using BI16, only CR[32–35] (bits within CR0) may be specified.

In branch conditional instructions, the BO32 or BO16 field specifies the conditions under which the branch is taken and how the branch is affected by or affects the CR and CTR. Note that VLE instructions also have different encodings for the BO32 and BO16 fields than in Book E's BO field.

If the BO32 field specifies that the CTR is to be decremented, CTR[32–63] are decremented. If BO[16,32] specifies a condition that must be TRUE or FALSE, that condition is obtained from the contents of CR[Bi+32]. (Note that CR bits are numbered 32–63. BI refers to the BI field in the branch instruction encoding. For example, specifying BI = 2 refers to CR[34].)

Encodings for the BO32 field for the VLE extension are shown in [Table 120](#).

**Table 119. VLE extension BO32 encodings**

BO32	Description
00	Branch if the condition is FALSE.
01	Branch if the condition is TRUE.
10	Decrement CTR[32–63], then branch if the decremented CTR[32–63]≠0.
11	Decrement CTR[32–63], then branch if the decremented CTR[32–63] = 0.

The encoding for the BO16 field for the VLE extension is shown in [Table 120](#).

**Table 120. VLE extension BO16 encodings**

BO16	Description
0	Branch if the condition is FALSE.
1	Branch if the condition is TRUE.

The various branch instructions supported by the VLE extension are shown in [Table 121](#).

Table 121. Branch instruction set index

Mnemonic	Instruction	Reference
<b>e_b</b> BD24 <b>e_bl</b> BD24	Branch Branch & Link	<a href="#">Section on page 744</a>
<b>se_b</b> BD8 <b>se_bl</b> BD8	Branch Branch & Link	<a href="#">Section on page 744</a>
<b>e_bc</b> BO32,BI32,BD15 <b>se_bc</b> BO16,BI16,BD8 <b>e_bcl</b> BO32,BI32,BD15	Branch Conditional Branch Conditional Branch Conditional & Link	<a href="#">Section on page 745</a>
<b>se_bctr</b> <b>se_bctrl</b>	Branch to Count Register Branch to Count Register & Link	<a href="#">Section on page 746</a>
<b>se_blr</b> <b>se_blrl</b>	Branch to Link Register Branch to Link Register & Link	<a href="#">Section on page 747</a>

#### 4.6.6.4 Condition register instructions

Condition register instructions are provided to transfer values to/from various portions of the CR. The VLE extension does not introduce any additional functionality beyond that defined in Book E for CR operations, but does remap the CR-logical and **mcrf** instruction functionality into major opcode 31. These instructions operate identically to the Book E instructions, but are encoded differently. [Table 122](#) lists condition register instructions supported in VLE mode.

Table 122. Condition register instruction set index

Mnemonic	Instruction	Reference
<b>e_crand</b> crbD,crbA,crbB	Condition Register AND	<a href="#">Section on page 753</a>
<b>e_crandc</b> crbD,crbA,crbB	Condition Register AND with Complement	<a href="#">Section on page 753</a>
<b>e_creqv</b> crbD,crbA,crbB	Condition Register Equivalent	<a href="#">Section on page 753</a>
<b>e_crnand</b> crbD,crbA,crbB	Condition Register NAND	<a href="#">Section on page 754</a>
<b>e_crnor</b> crbD,crbA,crbB	Condition Register NOR	<a href="#">Section on page 754</a>
<b>e_cror</b> crbD,crbA,crbB	Condition Register OR	<a href="#">Section on page 755</a>
<b>e_crorc</b> crbD,crbA,crbB	Condition Register OR with Complement	<a href="#">Section on page 755</a>
<b>e_crxor</b> crbD,crbA,crbB	Condition Register XOR	<a href="#">Section on page 755</a>



Table 122. Condition register instruction set index

Mnemonic	Instruction	Reference
<b>e_mcrf</b> crD,crS	Move Condition Register Field	<a href="#">Section on page 761</a>
<b>mcrxr</b> crD	Move to Condition Register from Integer Exception Register	Book E
<b>mfcrr</b> rD	Move From condition register	Book E
<b>mtcrf</b> FXM,rS	Move to Condition Register Fields	Book E

#### 4.6.6.5 Integer instructions

This section lists the integer instructions supported by the VLE extension.

##### 4.6.6.5.1 Integer load instructions

The integer load instructions, listed in [Table 123](#), compute the EA of the memory to be accessed as described in [Section 11.1: Data memory addressing modes](#).

The byte, half word, or word in memory addressed by EA is loaded into GPR(rD) or GPR(rZ).

The VLE extension supports both big- and little-endian byte ordering for data accesses.

Some integer load instructions have an update form in which GPR(rA) is updated with the EA. For these forms, if  $rA \neq 0$  and  $rA \neq rD$ , the EA is placed into GPR(rA) and the memory element (byte, half word, word, or double word) addressed by EA is loaded into GPR(rD). If  $rA = 0$  or  $rA = rD$ , the instruction form is invalid. This is the same behavior as specified for load with update instructions in Book E.

Table 123. Basic integer load instruction set index

Mnemonic	Instruction	Reference
<b>e_lbz</b> rD,D(rA) <b>e_lbzu</b> rD,D8(rA) <b>se_lbz</b> rZ,SD4(rX)	Load Byte and Zero Load Byte and Zero with Update Load Byte and Zero (16-bit form)	<a href="#">Section on page 757</a>
<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	Book E
<b>e_lha</b> rD,D(rA) <b>e_lhau</b> rD,D8(rA)	Load Halfword Algebraic Load Halfword Algebraic with Update	<a href="#">Section on page 758</a>
<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	Book E
<b>e_lhz</b> rD,D(rA) <b>e_lhzu</b> rD,D8(rA) <b>se_lhz</b> rZ,SD4(rX)	Load Halfword and Zero Load Halfword and Zero with Update Load Halfword and Zero (16-bit form)	<a href="#">Section on page 759</a>
<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	Book E

**Table 123. Basic integer load instruction set index**

Mnemonic	Instruction	Reference
<b>e_lwz</b> rD,D(rA)	Load Word and Zero	<a href="#">Section on page 761</a>
<b>e_lwzu</b> rD,D8(rA)	Load Word and Zero with Update	
<b>se_lwz</b> rZ,SD4(rX)	Load Word and Zero (16-bit form)	
<b>lwzx</b> rD,rA,rB	Load Word and Zero Indexed	Book E
<b>lwzux</b> rD,rA,rB	Load Word and Zero with Update Indexed	

Integer load byte-reversed instructions are listed in [Table 124](#).

**Table 124. Integer load byte-reverse instruction set index**

Mnemonic	Instruction	Reference
<b>lhbrx</b> rD,rA,rB	Load Halfword Byte-Reverse Indexed	Book E
<b>lwbrx</b> rD,rA,rB	Load Word Byte-Reverse Indexed	Book E

The VLE-defined integer load multiple instruction is listed in [Table 125](#).

**Table 125. Integer load multiple instruction set index**

Mnemonic	Instruction	Reference
<b>e_lmw</b> rD,D8(rA)	Load Multiple Word	<a href="#">Section on page 760</a>

The VLE-defined integer load and reserve instruction is listed in [Table 126](#).

**Table 126. Integer load and reserve instruction set index**

Mnemonic	Instruction	Reference
<b>lwarx</b> rD,rA,rB	Load Word And Reserve Indexed	Book E

#### 4.6.6.5.2 Integer store instructions

The integer store instructions compute the EA of the memory to be accessed as described in [Section 11.1: Data memory addressing modes.](#)

The contents of GPR(rS) or GPR(rZ) are stored into the byte, half word, or word in memory addressed by EA.

The VLE extension supports both big- and little-endian byte ordering for data accesses.

Some integer store instructions have an update form, in which GPR(rA) is updated with the EA. For these forms, the following rules (from Book E) apply.

- If  $rA \neq 0$ , the EA is placed into GPR(rA).
- If  $rS = rA$ , the contents of GPR(rS) are copied to the target memory element and then EA is placed into GPR(rA).

The basic integer store instructions are listed in [Table 127](#).

**Table 127. Basic integer store instruction set index**

Mnemonic	Instruction	Reference
<b>e_stb</b> rS,D(rA) <b>e_stbu</b> rS,D8(rA) <b>se_stb</b> rZ,SD4(rX)	Store Byte Store Byte with Update Store Byte (16-bit form)	<a href="#">Section on page 773</a>
<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Store Byte Indexed Store Byte with Update Indexed	Book E
<b>e_sth</b> rS,D(rA) <b>e_sthu</b> rS,D8(rA) <b>se_sth</b> rZ,SD4(rX)	Store Halfword Store Halfword with Update Store Halfword (16-bit form)	<a href="#">Section on page 773</a>
<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Store Halfword Indexed Store Halfword with Update Indexed	Book E
<b>e_stw</b> rS,D(rA) <b>e_stwu</b> rS,D8(rA) <b>se_stw</b> rZ,SD4(rX)	Store Word Store Word with Update Store Word (16-bit form)	<a href="#">Section on page 775</a>
<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Store Word Indexed Store Word with Update Indexed	Book E

The integer store byte-reverse instructions are listed in [Table 128](#).

**Table 128. Integer store byte-reverse instruction set index**

Mnemonic	Instruction	Reference
<b>sthrbx</b> rS,rA,rB	Store Halfword Byte-Reverse Indexed	Book E
<b>stwbrx</b> rS,rA,rB	Store Word Byte-Reverse Indexed	Book E

The integer store multiple instruction is listed in [Table 129](#).

**Table 129. Integer store multiple instruction set index**

Mnemonic	Instruction	Reference
<b>e_stmw</b> rS,D8(rA)	Store Multiple Word	<a href="#">Section on page 774</a>

The integer store conditional instruction is listed in [Table 130](#).

**Table 130. Integer store conditional instruction set index**

Mnemonic	Instruction	Reference
<b>stwcx.</b> rS,rA,rB	Store Word Conditional Indexed	Book E

#### 4.6.6.5.3 Integer arithmetic instructions

The integer arithmetic instructions use the contents of the GPRs as source operands, and place results into GPRs, into status bits in the XER and into CR0.

The integer arithmetic instructions treat source operands as signed, two's complement integers unless the instruction is explicitly identified as performing an unsigned operation.

The **e\_add2i**. instruction and the OIM5-form instruction, **se\_subi**., set the first three bits of CR0 to characterize bits 32–63 of the result. These bits are set by signed comparison of bits 32–63 of the result to zero.

**e\_addic**[.] and **e\_subfic**[.] always set CA to reflect the carry out of bit 32.

The integer arithmetic instructions are listed in [Table 131](#).

**Table 131. Integer arithmetic instruction set index**

Mnemonic	Instruction	Reference
<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Add	Book E
<b>se_add</b> rX,rY	Add	<a href="#">Section on page 741</a>
<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Add Carrying	Book E
<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Add Extended	Book E
<b>e_addi</b> rD,rA,SC18 <b>e_addi.</b> rD,rA,SC18 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI <b>se_addi</b> rX,OIMM	Add Immediate	<a href="#">Section on page 741</a>
<b>e_addic</b> rD,rA,SC18 <b>e_addic.</b> rD,rA,SC18	Add Immediate Carrying	<a href="#">Section on page 743</a>
<b>e_add2is</b> rD,SI	Add Immediate Shifted	<a href="#">Section on page 741</a>
<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Divide Word	Book E
<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Divide Word Unsigned	Book E
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word	Book E

Table 131. Integer arithmetic instruction set index (continued)

Mnemonic	Instruction	Reference
<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Multiply High Word Unsigned	Book E
<b>e_mulli</b> rD,rA,SCI8 <b>e_mull2i</b> rD,SI	Multiply Low Immediate	<a href="#">Section on page 763</a>
<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Multiply Low Word	Book E
<b>se_mullw</b> rX,rY	Multiply Low Word	<a href="#">Section on page 764</a>
<b>neg</b> rD,rA <b>se_neg</b> rX <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Negate	<a href="#">Section on page 764</a>
<b>se_sub</b> rX,rY	Subtract	<a href="#">Section on page 775</a>
<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Subtract From	Book E
<b>se_subf</b> rX,rY	Subtract From	<a href="#">Section on page 776</a>
<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Subtract From Carrying	Book E
<b>e_subfic</b> rD,rA,SCI8 <b>e_subfic.</b> rD,rA,SCI8	Subtract From Immediate Carrying	<a href="#">Section on page 776</a>
<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	Subtract Immediate	<a href="#">Section on page 776</a>

#### 4.6.6.5.4 Integer logical and move instructions

Logical instructions perform bit-parallel operations on 32-bit operands or move register or immediate values into registers. The move instructions move values into a GP from either another GPR, or an immediate value.

The X-form logical instructions with Rc = 1 and the SCI8-form logical instructions with Rc = 1 set the first three bits of CR field 0 as described in [Section 4.6.6.5.3: Integer arithmetic instructions](#). The logical instructions do not change XER[SO,OV,CA].

The integer logical instructions are listed in [Table 132](#).

Table 132. Integer logical instruction set index

Mnemonic	Instruction	Reference
<b>and</b> [.] rA,rS,rB <b>se_and</b> [.] rX,rY	AND	<a href="#">Section on page 743</a>
<b>andc</b> [.] rA,rS,rB <b>se_andc</b> rX,rY	AND with Complement	<a href="#">Section on page 743</a>
<b>e_andi</b> [.] rA,rS,SCI8 <b>se_andi</b> rX,UI5 <b>e_and2i</b> . rD,UI	AND Immediate	<a href="#">Section on page 743</a>
<b>e_and2is</b> . rD,UI	AND Immediate Shifted	<a href="#">Section on page 743</a>
<b>se_bclri</b> rX,UI5	Bit Clear	<a href="#">Section on page 746</a>
<b>se_bgeni</b> rX,UI5	Bit Generate	<a href="#">Section on page 747</a>
<b>se_bmski</b> rX,UI5	Bit Mask Generate	<a href="#">Section on page 748</a>
<b>se_bseti</b> rX,UI5	Bit Set	<a href="#">Section on page 748</a>
<b>cntlzw</b> rA,rS <b>cntlzw</b> . rA,rS	Count Leading Zeros Word	Book E
<b>eqv</b> rA,rS,rB <b>eqv</b> . rA,rS,rB	Equivalent	Book E
<b>extsb</b> rA,rS <b>extsb</b> . rA,rS <b>se_extsb</b> rX	Extend Sign Byte	<a href="#">Section on page 755</a>
<b>extsh</b> rA,rS <b>extsh</b> . rA,rS <b>se_extsh</b> rX	Extend Sign Halfword	<a href="#">Section on page 755</a>
<b>se_extzb</b> rX	Extend with Zeros Byte	<a href="#">Section on page 756</a>
<b>se_extzh</b> rX	Extend with Zeros Halfword	<a href="#">Section on page 756</a>
<b>e_li</b> rD,LI20 <b>se_li</b> rX,UI7	Load Immediate	<a href="#">Section on page 760</a>
<b>e_lis</b> rD,UI	Load Immediate Shifted	<a href="#">Section on page 760</a>
<b>se_mfar</b> rX,arY	Move from Alternate Register	<a href="#">Section on page 762</a>
<b>se_mr</b> rX,rY	Move Register	<a href="#">Section on page 762</a>
<b>se_mtar</b> arX,rY	Move to Alternate Register	<a href="#">Section on page 763</a>

Table 132. Integer logical instruction set index (continued)

Mnemonic	Instruction	Reference
<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	NAND	Book E
<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	NOR	Book E
<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB <b>se_or</b> rX,rY	OR	<a href="#">Section on page 765</a>
<b>se_not</b> rX	NOT	<a href="#">Section on page 764</a>
<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	OR with Complement	Book E
<b>e_ori</b> [.] rA,rS,SCI8 <b>e_or2i</b> rD,UI	OR Immediate	<a href="#">Section on page 765</a>
<b>e_or2is</b> rD,UI	OR Immediate Shifted	<a href="#">Section on page 765</a>
<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	XOR	Book E
<b>e_xori</b> [.] rA,rS,SCI8	XOR Immediate	<a href="#">Section on page 765</a>

#### 4.6.6.5.5 Integer compare and bit test instructions

The integer compare instructions compare the contents of GPR(rA) with one of the following:

- The value of the SCI8 field
- The zero-extended value of the UI field
- The zero-extended value of the UI5 field
- The sign-extended value of the SI field
- The contents of GPR(rB) or GPR(rY).

The following comparisons are signed: **e\_cmph**, **e\_cmpi**, **e\_cmp16i**, **e\_cmph16i**, **se\_cmp**, **se\_cmph**, and **se\_cmpi**.

The following comparisons are unsigned: **e\_cmphl**, **e\_cmpli**, **e\_cmphl16i**, **e\_cmpl16i**, **se\_cmpli**, **se\_cmpl**, and **se\_cmphl**.

When operands are treated as 32-bit signed quantities, GPR $n$ [32] is the sign bit. When operands are treated as 16-bit signed quantities, GPR $n$ [48] is the sign bit.

For 32-bit implementations, the L field must be zero.

Compare instructions set one of the left-most three bits of the designated CR field and clears the other two. XER[SO] is copied to bit 3 of the designated CR field.

The CR field is set as shown in [Table 133](#).

**Table 133. CR settings for compare instructions**

Bit	Name	Description
0	LT	(rA or rX) < SCI8, SI, UI5, or GPR(rB or rY) (signed comparison) (rA or rX) < <sub>u</sub> SCI8, UI, UI5 or GPR(rB or rY) (unsigned comparison)
1	GT	(rA or rX) > SCI8, SI, UI5, or GPR(rB or rY) (signed comparison) (rA or rX) > <sub>u</sub> SCI8, UI, UI5 or GPR(rB or rY) (unsigned comparison)
2	EQ	(rA or rX) = SCI8, SI, UI, UI5, or GPR(rB or rY)
3	SO	Summary overflow from the XER

The integer bit test instruction tests the bit specified by the UI5 instruction field and sets the CR0 field as shown in [Table 134](#).

**Table 134. CR settings for integer bit test instructions**

Bit	Name	Description
0	LT	Always cleared
1	GT	RX <sub>ui5</sub> == 1
2	EQ	RX <sub>ui5</sub> == 0
3	SO	Summary overflow from the XER

[Table 135](#) is an index for integer compare and bit test operations.

**Table 135. Integer compare and bit test instruction set index**

Mnemonic	Instruction	Reference
<b>se_btsti</b> rX,UI5	Bit Test Immediate	<a href="#">Section on page 748</a>
<b>cmp</b> crD,L,rA,rB <b>se_cmp</b> rX,rY	Compare	<a href="#">Section on page 749</a>
<b>e_cmph</b> crD,rA,rB <b>se_cmph</b> rX,rY	Compare Halfword	<a href="#">Section on page 750</a>
<b>e_cmph16i</b> rA,SI16	Compare Halfword Immediate	<a href="#">Section on page 750</a>
<b>e_cmphl</b> crD,rA,rB <b>se_cmphl</b> rX,rY	Compare Halfword Logical	<a href="#">Section on page 751</a>
<b>e_cmphl16i</b> rA,UI16	Compare Halfword Logical Immediate	<a href="#">Section on page 751</a>
<b>e_cmpi</b> crD,rA,SCI8 <b>e_cmp16i</b> rA,SI16 <b>se_cmpi</b> rX,UI5	Compare Immediate	<a href="#">Section on page 749</a>



**Table 135. Integer compare and bit test instruction set index (continued)**

Mnemonic	Instruction	Reference
<b>cmpl</b> crD,L,rA,rB <b>se_cmpl</b> rX,rY	Compare Logical	<a href="#">Section on page 752</a>
<b>e_cmpli</b> crD,rA,SCI8 <b>e_cmpl16i</b> rA,UI16 <b>se_cmpli</b> rX,UI5	Compare Logical Immediate	<a href="#">Section on page 752</a>

#### 4.6.6.5.6 Integer select instruction

The **isel** instruction provides a means to select one of two registers and place the result in a destination register under the control of a predicate value supplied by a CR bit.

The integer select instruction is listed in [Table 136](#).

**Table 136. Integer select instruction set index**

Mnemonic	Instruction	Reference
<b>isel</b> rD,rA,rB,crb	Integer Select	EIS

#### 4.6.6.5.7 Integer trap instructions

Trap instructions test for a specified set of conditions by comparing the contents of one GPR with a second GPR. If any of the conditions tested by a Trap instruction are met, a trap exception type program interrupt is invoked. If none of the tested conditions are met, instruction execution continues normally.

The contents of GPR(rA) are compared with the contents of GPR(rB). For **twi** and **tw**, only the contents of bits 32–63 of rA (and rB) participate in the comparison.

This comparison results in five conditions that are ANDed with TO. If the result is not 0, the trap exception type program interrupt is invoked. These conditions are as shown in [Table 137](#).

**Table 137. Integer trap conditions**

TO Bit	ANDed with condition
0	Less Than, using signed comparison
1	Greater Than, using signed comparison
2	Equal
3	Less Than, using unsigned comparison
4	Greater Than, using unsigned comparison

The integer trap instruction is listed in [Table 138](#).

**Table 138. Integer trap instruction set index**

Mnemonic	Instruction	Reference
<b>tw</b> TO,rA,rB	Trap Word	Book E

**4.6.6.5.8 Integer rotate and shift instructions**

Instructions are provided that perform shifts and rotates on data from a GPR and return the result, or a portion of the result, to a GPR.

The rotation operations rotate a 32-bit quantity left by a specified number of bit positions. Bits that exit from position 32 enter at position 63.

The rotate<sub>32</sub> operation is used to rotate a given 32-bit quantity.

Some rotate and shift instructions employ a mask generator. The mask is 32 bits long, and consists of 1 bits from a start bit, *mstart*, through and including a stop bit, *mstop*, and 0-bits elsewhere. The values of *mstart* and *mstop* range from 32 to 63. If *mstart* > *mstop*, the 1 bits wrap around from position 63 to position 0. Thus the mask is formed as follows:

```

if mstart ≤ mstop then
    maskmstart:mstop = ones
    maskall other bits = zeros
else
    maskmstart:63 = ones
    mask32:mstop = ones
    maskall other bits = zeros
    
```

There is no way to specify an all-zero mask.

For instructions that use the rotate<sub>32</sub> operation, the mask start and stop positions are always in bits 32–63 of the mask.

The use of the mask is described in following sections.

The rotate word and shift word instructions with Rc = 1 set the first three bits of CR field 0 as described in Book E. Rotate and shift instructions do not change the OV and SO bits. Rotate and shift instructions, except algebraic right shifts, do not change the CA bit.

The instructions in [Table 139](#) rotate the contents of a register. Depending on the instruction type, the amount of the rotation is either specified as an immediate, or contained in a GPR.

**Table 139. Integer rotate instruction set index**

Mnemonic	Instruction	Reference
<b>e_rlw</b> rA,rS,rB	Rotate Left Word	<a href="#">Section on page 768</a>
<b>e_rlwi</b> rA,rS,SH	Rotate Left Word Immediate	<a href="#">Section on page 768</a>

The instructions in [Table 140](#) rotate the contents of a register. Depending on the instruction type, the result of the rotation is either inserted into the target register under control of a

mask (if a mask bit is 1, the associated bit of the rotated data is placed into the target register; if a mask bit is 0, the associated bit in the target register remains unchanged) or ANDed with a mask before being placed into the target register.

The rotate left instructions allow right-rotation of the contents of a register to be performed (in concept) by a left-rotation of  $32-n$ , where  $n$  is the number of bits by which to rotate right. They allow right-rotation of the contents of bits 32–63 of a register to be performed (in concept) by a left-rotation of  $32-n$ , where  $n$  is the number of bits by which to rotate right.

**Table 140. Integer rotate with mask instruction set index**

Mnemonic	Instruction	Reference
<b>e_rlwimi</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then Mask Insert	<a href="#">Section on page 768</a>
<b>e_rlwinm</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then AND with Mask	<a href="#">Section on page 769</a>

The integer shift instructions are listed in [Table 141](#).

**Table 141. Integer shift instruction set index**

Mnemonic	Instruction	Reference
<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB <b>se_slw</b> rX,rY	Shift Left Word	<a href="#">Section on page 770</a>
<b>e_slwi</b> rA,rS,SH <b>se_slwi</b> rX,UI5	Shift Left Word Immediate	<a href="#">Section on page 770</a>
<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB <b>se_sraw</b> rX,rY	Shift Right Algebraic Word	<a href="#">Section on page 771</a>
<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH <b>se_srawi</b> rX,UI5	Shift Right Algebraic Word Immediate	<a href="#">Section on page 771</a>
<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB <b>se_srw</b> rX,rY	Shift Right Word	<a href="#">Section on page 772</a>
<b>e_srwi</b> rA,rS,SH <b>se_srwi</b> rX,UI5	Shift Right Word Immediate	<a href="#">Section on page 772</a>

#### 4.6.6.6 Storage control instructions

This section lists storage control instructions, which include the following:

- [Section 4.6.6.7: Storage synchronization instructions](#)
- [Section 4.6.6.8: Cache management instructions](#)
- [Section 4.6.6.9: TLb management instructions](#)

#### 4.6.6.7 Storage synchronization instructions

The memory synchronization instructions implemented by the VLE extension are identical to those defined in Book E.

The storage synchronization instructions are listed in [Table 142](#).

**Table 142. Storage synchronization instruction set index**

Mnemonic	Instruction	Reference
<b>mbar</b>	Memory Barrier	Book E
<b>msync</b>	Memory Synchronize	Book E

#### 4.6.6.8 Cache management instructions

Cache management instructions implemented by the VLE extension are identical to those defined in Book E.

The cache management instructions are listed in [Table 143](#).

**Table 143. Cache management instruction set index**

Mnemonic	Instruction	Reference
<b>dcba rA,rB</b>	Data Cache Block Allocate	Book E
<b>dcbf rA,rB</b>	Data Cache Block Flush	Book E
<b>dcbi rA,rB</b>	Data Cache Block Invalidate	Book E
<b>dcbst rA,rB</b>	Data Cache Block Store	Book E
<b>dcbt CT,rA,rB</b>	Data Cache Block Touch	Book E
<b>dcbtst CT,rA,rB</b>	Data Cache Block Touch for Store	Book E
<b>dcbz rA,rB</b>	Data Cache Block set to Zero	Book E
<b>icbi rA,rB</b>	Instruction Cache Block Invalidate	Book E
<b>icbt CT,rA,rB</b>	Instruction Cache Block Touch	Book E

#### 4.6.6.9 TLB management instructions

The TLB management instructions implemented by the VLE extension are identical to those defined in Book E and in the EIS. The TLB management instructions are listed in [Table 144](#).

**Table 144. TLB management instruction set index**

Mnemonic	Instruction	Reference
<b>tlbivax rA,rB</b>	TLB Invalidate Virtual Address Indexed	Book E
<b>tlbre</b>	TLB Read Entry	Book E
<b>tlbsx rA,rB</b>	TLB Search Indexed	Book E
<b>tlbsync</b>	TLB Synchronize	Book E
<b>tlbwe</b>	TLB Write Entry	Book E

#### 4.6.6.10 VLE instruction alignment and byte ordering

An instruction fetched from memory must be placed in the pipeline with its bytes in the proper order. Otherwise, the instruction decoder cannot recognize it. Book E allows instructions to be placed into memory marked as either big- or little-endian. This is manageable because Book E instructions are always word-size aligned on word boundaries. The VLE extension includes both half-word- and word-length instructions are aligned on half-word boundaries. Because of this, only big-endian instruction memory is supported when executing from a page of VLE instructions. Attempts to execute VLE instructions from a page marked as little-endian generate an instruction storage interrupt byte-ordering exception.

#### 4.6.6.11 Instruction listings

This section lists instructions either defined or supported by the VLE extension.

[Table 145](#) lists instructions by instruction name.

**Table 145. Instructions listed by name**

Instruction	Mnemonic	Reference
Add	<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Book E
Add Carrying	<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Book E
Add Extended	<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Book E
AND with Complement	<b>andc</b> [.] rA,rS,rB <b>se_andc</b> rX,rY	Book E <a href="#">Section on page 743</a>
AND	<b>and</b> [.] rA,rS,rB <b>se_and</b> [.] rX,rY	Book E <a href="#">Section on page 743</a>
Compare	<b>cmp</b> crD,L,rA,rB <b>se_cmp</b> rX,rY	Book E <a href="#">Section on page 749</a>
Compare Logical	<b>cmpl</b> crD,L,rA,rB <b>se_cmpl</b> rX,rY	Book E <a href="#">Section on page 752</a>
Count Leading Zeros Word	<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Book E
Data Cache Block Allocate	<b>dcba</b> rA,rB	Book E
Data Cache Block Flush	<b>dcbf</b> rA,rB	Book E

**Table 145. Instructions listed by name (continued)**

Instruction	Mnemonic	Reference
Data Cache Block Invalidate	<b>dcbi</b> rA,rB	Book E
Data Cache Block Store	<b>dcbst</b> rA,rB	Book E
Data Cache Block Touch	<b>dcbt</b> CT,rA,rB	Book E
Data Cache Block Touch for Store	<b>dcbtst</b> CT,rA,rB	Book E
Data Cache Block set to Zero	<b>dcbz</b> rA,rB	Book E
Divide Word	<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Book E
Divide Word Unsigned	<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Book E
Equivalent	<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Book E
Extend Sign Byte	<b>extsb</b> rA,rS <b>extsb.</b> rA,rS <b>se_extsb</b> rX	Book E Book E <a href="#">Section on page 755</a>
Extend Sign Halfword	<b>extsh</b> rA,rS <b>extsh.</b> rA,rS <b>se_extsh</b> rX	Book E Book E <a href="#">Section on page 755</a>
Add Immediate Shifted	<b>e_add2is</b> rD,SI	<a href="#">Section on page 741</a>
Add Immediate	<b>e_addi</b> rD,rA,SCI8 <b>e_addi.</b> rD,rA,SCI8 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI <b>se_addi</b> rX,OIMM	<a href="#">Section on page 741</a>
Add Immediate Carrying	<b>e_addic</b> rD,rA,SCI8 <b>e_addic.</b> rD,rA,SCI8	<a href="#">Section on page 743</a>
AND Immediate Shifted	<b>e_and2is.</b> rD,UI	<a href="#">Section on page 743</a>
AND Immediate	<b>e_andi[.]</b> rA,rS,SCI8 <b>se_andi</b> rX,UI5 <b>e_and2i.</b> rD,UI	<a href="#">Section on page 743</a>
Branch Conditional Branch Conditional Branch Conditional & Link	<b>e_bc</b> BO32,BI32,BD15 <b>se_bc</b> BO16,BI16,BD8 <b>e_bcl</b> BO32,BI32,BD15	<a href="#">Section on page 745</a>

Table 145. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Branch	<b>e_b</b> BD24	<a href="#">Section on page 744</a>
Branch & Link	<b>e_bl</b> BD24	<a href="#">Section on page 744</a>
Compare Halfword	<b>e_cmp</b> crD,rA,rB <b>se_cmp</b> rX,rY	<a href="#">Section on page 750</a>
Compare Halfword Immediate	<b>e_cmp</b> 16i rA,SI16	<a href="#">Section on page 750</a>
Compare Halfword Logical	<b>e_cmpl</b> crD,rA,rB <b>se_cmpl</b> rX,rY	<a href="#">Section on page 751</a>
Compare Halfword Logical Immediate	<b>e_cmpl</b> 16i rA,UI16	<a href="#">Section on page 751</a>
Compare Immediate	<b>e_cm</b> pi crD,rA,SCI8 <b>e_cm</b> 16i rA,SI16 <b>se_cm</b> pi rX,UI5	<a href="#">Section on page 749</a>
Compare Logical Immediate	<b>e_cm</b> pli crD,rA,SCI8 <b>e_cm</b> pl16i rA,UI16 <b>se_cm</b> pli rX,UI5	<a href="#">Section on page 752</a>
Condition Register AND	<b>e_c</b> rand crbD,crbA,crbB	<a href="#">Section on page 753</a>
Condition Register AND with Complement	<b>e_c</b> randc crbD,crbA,crbB	<a href="#">Section on page 753</a>
Condition Register Equivalent	<b>e_c</b> reqv crbD,crbA,crbB	<a href="#">Section on page 753</a>
Condition Register NAND	<b>e_c</b> rnand crbD,crbA,crbB	<a href="#">Section on page 754</a>
Condition Register NOR	<b>e_c</b> rnor crbD,crbA,crbB	<a href="#">Section on page 754</a>
Condition Register OR	<b>e_c</b> ror crbD,crbA,crbB	<a href="#">Section on page 755</a>
Condition Register OR with Complement	<b>e_c</b> rorc crbD,crbA,crbB	<a href="#">Section on page 755</a>
Condition Register XOR	<b>e_c</b> rxor crbD,crbA,crbB	<a href="#">Section on page 755</a>
Load Byte and Zero	<b>e_l</b> bz rD,D(rA)	<a href="#">Section on page 757</a>
Load Byte and Zero with Update	<b>e_l</b> bzu rD,D8(rA)	
Load Byte and Zero (16-bit form)	<b>se_l</b> bz rZ,SD4(rX)	
Load Halfword Algebraic	<b>e_l</b> ha rD,D(rA)	<a href="#">Section on page 758</a>
Load Halfword Algebraic with Update	<b>e_l</b> hau rD,D8(rA)	
Load Halfword and Zero	<b>e_l</b> hz rD,D(rA)	<a href="#">Section on page 759</a>
Load Halfword and Zero with Update	<b>e_l</b> hzu rD,D8(rA)	
Load Halfword and Zero (16-bit form)	<b>se_l</b> hz rZ,SD4(rX)	
Load Immediate	<b>e_l</b> i rD,LI20 <b>se_l</b> i rX,UI7	<a href="#">Section on page 760</a>

Table 145. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Load Immediate Shifted	<b>e_lis</b> rD,UI	<a href="#">Section on page 760</a>
Load Multiple Word	<b>e_lmw</b> rD,D8(rA)	<a href="#">Section on page 760</a>
Load Word and Zero Load Word and Zero with Update Load Word and Zero (16-bit form)	<b>e_lwz</b> rD,D(rA) <b>e_lwzu</b> rD,D8(rA) <b>se_lwz</b> rZ,SD4(rX)	<a href="#">Section on page 761</a>
Move Condition Register Field	<b>e_mcrf</b> crD,crS	<a href="#">Section on page 761</a>
Multiply Low Immediate	<b>e_mulli</b> rD,rA,SCI8 <b>e_mull2i</b> rD,SI	<a href="#">Section on page 763</a>
OR Immediate Shifted	<b>e_or2is</b> rD,UI	<a href="#">Section on page 765</a>
OR Immediate	<b>e_ori</b> [.] rA,rS,SCI8 <b>e_or2i</b> rD,UI	<a href="#">Section on page 765</a>
Rotate Left Word	<b>e_rlw</b> rA,rS,rB	<a href="#">Section on page 768</a>
Rotate Left Word Immediate	<b>e_rlwi</b> rA,rS,SH	<a href="#">Section on page 768</a>
Rotate Left Word Immediate then Mask Insert	<b>e_rlwimi</b> rA,rS,SH,MB,ME	<a href="#">Section on page 768</a>
Rotate Left Word Immediate then AND with Mask	<b>e_rlwinm</b> rA,rS,SH,MB,ME	<a href="#">Section on page 769</a>
Shift Left Word Immediate	<b>e_slwi</b> rA,rS,SH <b>se_slwi</b> rX,UI5	<a href="#">Section on page 770</a>
Shift Right Word Immediate	<b>e_srwi</b> rA,rS,SH <b>se_srwi</b> rX,UI5	<a href="#">Section on page 772</a>
Store Byte Store Byte with Update Store Byte (16-bit form)	<b>e_stb</b> rS,D(rA) <b>e_stbu</b> rS,D8(rA) <b>se_stb</b> rZ,SD4(rX)	<a href="#">Section on page 773</a>
Store Halfword Store Halfword with Update Store Halfword (16-bit form)	<b>e_sth</b> rS,D(rA) <b>e_sthu</b> rS,D8(rA) <b>se_sth</b> rZ,SD4(rX)	<a href="#">Section on page 773</a>
Store Multiple Word	<b>e_stmw</b> rS,D8(rA)	<a href="#">Section on page 774</a>
Store Word Store Word with Update Store Word (16-bit form)	<b>e_stw</b> rS,D(rA) <b>e_stwu</b> rS,D8(rA) <b>se_stw</b> rZ,SD4(rX)	<a href="#">Section on page 775</a>
Subtract From Immediate Carrying	<b>e_subfic</b> rD,rA,SCI8 <b>e_subfic.</b> rD,rA,SCI8	<a href="#">Section on page 776</a>



Table 145. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
XOR Immediate	<b>e_xori</b> [.] rA,rS,SCI8	<a href="#">Section on page 765</a>
Instruction Cache Block Invalidate	<b>icbi</b> rA,rB	Book E
Instruction Cache Block Touch	<b>icbt</b> CT,rA,rB	Book E
Integer Select	<b>isel</b> rD,rA,rB,crb	EIS
Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Book E
Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Book E
Load Halfword Byte-Reverse Indexed	<b>lhbrx</b> rD,rA,rB	Book E
Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Book E
Load Word And Reserve Indexed	<b>lwarx</b> rD,rA,rB	Book E
Load Word Byte-Reverse Indexed	<b>lwbrx</b> rD,rA,rB	Book E
Load Word and Zero Indexed Load Word and Zero with Update Indexed	<b>lwzx</b> rD,rA,rB <b>lwzux</b> rD,rA,rB	Book E
Memory Barrier	<b>mbar</b>	Book E
Move to Condition Register from Integer Exception Register	<b>mcrxr</b> crD	Book E
Move From condition register	<b>mfcrr</b> rD	Book E
Move From Device Control Register	<b>mfdcr</b> rD,DCRN	Book E
Move From Machine State Register	<b>mfmsr</b> rD	Book E
Move From Special Purpose Register	<b>mfspir</b> rD,SPRN	Book E
Memory Synchronize	<b>msync</b>	Book E
Move to Condition Register Fields	<b>mtrcrf</b> FXM,rS	Book E
Move To Device Control Register	<b>mtdcr</b> DCRN,rS	Book E
Move To Machine State Register	<b>mtmsr</b> rS	Book E
Move To Special Purpose Register	<b>mtspr</b> SPRN,rS	Book E
Multiply High Word	<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Book E
Multiply High Word Unsigned	<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Book E
Multiply Low Word	<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Book E
NAND	<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	Book E

**Table 145. Instructions listed by name (continued)**

Instruction	Mnemonic	Reference
Negate	<b>neg</b> rD,rA <b>se_neg</b> rX <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Book E <a href="#">Section on page 764</a> Book E Book E Book E
NOR	<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	Book E
OR	<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB <b>se_or</b> rX,rY	Book E Book E <a href="#">Section on page 765</a>
OR with Complement	<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	Book E
Add	<b>se_add</b> rX,rY	<a href="#">Section on page 741</a>
Bit Clear	<b>se_bclr</b> rX,UI5	<a href="#">Section on page 746</a>
Branch to Count Register Branch to Count Register & Link	<b>se_bctr</b> <b>se_bctrl</b>	<a href="#">Section on page 746</a>
Bit Generate	<b>se_bgeni</b> rX,UI5	<a href="#">Section on page 747</a>
Branch to Link Register Branch to Link Register & Link	<b>se_blr</b> <b>se_blrl</b>	<a href="#">Section on page 747</a>
Bit Mask Generate	<b>se_bmski</b> rX,UI5	<a href="#">Section on page 748</a>
Bit Set	<b>se_bseti</b> rX,UI5	<a href="#">Section on page 748</a>
Branch Branch & Link	<b>se_b</b> BD8 <b>se_bl</b> BD8	<a href="#">Section on page 744</a>
Bit Test Immediate	<b>se_btsti</b> rX,UI5	<a href="#">Section on page 748</a>
Extend with Zeros Byte	<b>se_extzb</b> rX	<a href="#">Section on page 756</a>
Extend with Zeros Halfword	<b>se_extzh</b> rX	<a href="#">Section on page 756</a>
Instruction Synchronize	<b>se_isync</b>	<a href="#">Section on page 757</a>
Move from Alternate Register	<b>se_mfar</b> rX,arY	<a href="#">Section on page 762</a>
Move From Count Register	<b>se_mfctr</b> rX	<a href="#">Section on page 762</a>

Table 145. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Move From Link Register	<b>se_mflr</b> rX	<a href="#">Section on page 762</a>
Move Register	<b>se_mr</b> rX,rY	<a href="#">Section on page 762</a>
Move to Alternate Register	<b>se_mtar</b> arX,rY	<a href="#">Section on page 763</a>
Move To Count Register	<b>se_mtctr</b> rX	<a href="#">Section on page 763</a>
Move To Link Register	<b>se_mtlr</b> rX	<a href="#">Section on page 763</a>
Multiply Low Word	<b>se_mullw</b> rX,rY	<a href="#">Section on page 764</a>
NOT	<b>se_not</b> rX	<a href="#">Section on page 764</a>
Subtract	<b>se_sub</b> rX,rY	<a href="#">Section on page 775</a>
Subtract From	<b>se_subf</b> rX,rY	<a href="#">Section on page 776</a>
Subtract Immediate	<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	<a href="#">Section on page 776</a>
Shift Left Word	<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB <b>se_slw</b> rX,rY	Book E Book E <a href="#">Section on page 770</a>
Shift Right Algebraic Word	<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB <b>se_sraw</b> rX,rY	Book E Book E <a href="#">Section on page 771</a>
Shift Right Algebraic Word Immediate	<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH <b>se_srawi</b> rX,UI5	Book E Book E <a href="#">Section on page 771</a>
Shift Right Word	<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB <b>se_srw</b> rX,rY	Book E Book E <a href="#">Section on page 772</a>
Store Byte Indexed Store Byte with Update Indexed	<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Book E
Store Halfword Byte-Reverse Indexed	<b>sthbrx</b> rS,rA,rB	Book E
Store Halfword Indexed Store Halfword with Update Indexed	<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Book E

**Table 145. Instructions listed by name (continued)**

Instruction	Mnemonic	Reference
Store Word Byte-Reverse Indexed	<b>stwbrx</b> rS,rA,rB	Book E
Store Word Conditional Indexed	<b>stwcx.</b> rS,rA,rB	Book E
Store Word Indexed Store Word with Update Indexed	<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Book E
Subtract From	<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Book E
Subtract From Carrying	<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Book E
TLB Invalidate Virtual Address Indexed	<b>tlbivax</b> rA,rB	Book E
TLB Read Entry	<b>tlbre</b>	Book E
TLB Search Indexed	<b>tlbsx</b> rA,rB	Book E
TLB Synchronize	<b>tlbsync</b>	Book E
TLB Write Entry	<b>tlbwe</b>	Book E
Trap Word	<b>tw</b> TO,rA,rB	Book E
Write MSR External Enable	<b>wrtee</b> rA	Book E
Write MSR External Enable Immediate	<b>wrteei</b> E	Book E
XOR	<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	Book E

*Table 145* lists instructions that can be executed in VLE mode by mnemonic.

**Table 146. Instructions listed by mnemonic**

Mnemonic	Instruction	Reference
<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Add	Book E
<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Add Carrying	Book E
<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Add Extended	Book E

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>andc</b> [.] rA,rS,rB	AND with Complement	Book E
<b>and</b> [.] rA,rS,rB	AND	Book E
<b>cmp</b> crD,L,rA,rB	Compare	Book E
<b>cmpl</b> crD,L,rA,rB	Compare Logical	Book E
<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Count Leading Zeros Word	Book E
<b>dcba</b> rA,rB	Data Cache Block Allocate	Book E
<b>dcbf</b> rA,rB	Data Cache Block Flush	Book E
<b>dcbi</b> rA,rB	Data Cache Block Invalidate	Book E
<b>dcbst</b> rA,rB	Data Cache Block Store	Book E
<b>dcbt</b> CT,rA,rB	Data Cache Block Touch	Book E
<b>dcbtst</b> CT,rA,rB	Data Cache Block Touch for Store	Book E
<b>dcbz</b> rA,rB	Data Cache Block set to Zero	Book E
<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Divide Word	Book E
<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Divide Word Unsigned	Book E
<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Equivalent	Book E
<b>extsb</b> rA,rS <b>extsb.</b> rA,rS	Extend Sign Byte	Book E
<b>extsh</b> rA,rS <b>extsh.</b> rA,rS	Extend Sign Halfword	Book E
<b>e_add2is</b> rD,SI	Add Immediate Shifted	<a href="#">Section on page 741</a>
<b>e_addi</b> rD,rA,SCI8 <b>e_addi.</b> rD,rA,SCI8 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI	Add Immediate	<a href="#">Section on page 741</a>
<b>e_addic</b> rD,rA,SCI8 <b>e_addic.</b> rD,rA,SCI8	Add Immediate Carrying	<a href="#">Section on page 743</a>
<b>e_and2is.</b> rD,UI	AND Immediate Shifted	<a href="#">Section on page 743</a>
<b>e_andi</b> [.] rA,rS,SCI8 <b>e_and2i.</b> rD,UI	AND Immediate	<a href="#">Section on page 743</a>

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>e_bc</b> BO32,BI32,BD15 <b>e_bcl</b> BO32,BI32,BD15	Branch Conditional Branch Conditional & Link	<a href="#">Section on page 745</a> <a href="#">Page</a>
<b>e_b</b> BD24 <b>e_bl</b> BD24	Branch Branch & Link	<a href="#">Section on page 744</a>
<b>e_cmph</b> crD,rA,rB	Compare Halfword	<a href="#">Section on page 750</a>
<b>e_cmph16i</b> rA,S116	Compare Halfword Immediate	<a href="#">Section on page 750</a>
<b>e_cmphl</b> crD,rA,rB	Compare Halfword Logical	<a href="#">Section on page 751</a>
<b>e_cmphl16i</b> rA,UI16	Compare Halfword Logical Immediate	<a href="#">Section on page 751</a>
<b>e_cmpi</b> crD,rA,SCI8 <b>e_cmp16i</b> rA,S116	Compare Immediate	<a href="#">Section on page 749</a>
<b>e_cmpli</b> crD,rA,SCI8 <b>e_cmpl16i</b> rA,UI16	Compare Logical Immediate	<a href="#">Section on page 752</a>
<b>e_crand</b> crbD,crbA,crbB	Condition Register AND	<a href="#">Section on page 753</a>
<b>e_crandc</b> crbD,crbA,crbB	Condition Register AND with Complement	<a href="#">Section on page 753</a>
<b>e_creqv</b> crbD,crbA,crbB	Condition Register Equivalent	<a href="#">Section on page 753</a>
<b>e_crnand</b> crbD,crbA,crbB	Condition Register NAND	<a href="#">Section on page 754</a>
<b>e_crnor</b> crbD,crbA,crbB	Condition Register NOR	<a href="#">Section on page 754</a>
<b>e_cror</b> crbD,crbA,crbB	Condition Register OR	<a href="#">Section on page 755</a>
<b>e_crorc</b> crbD,crbA,crbB	Condition Register OR with Complement	<a href="#">Section on page 755</a>
<b>e_crxor</b> crbD,crbA,crbB	Condition Register XOR	<a href="#">Section on page 755</a>
<b>e_lbz</b> rD,D(rA) <b>e_lbzu</b> rD,D8(rA)	Load Byte and Zero Load Byte and Zero with Update	<a href="#">Section on page 757</a>
<b>e_lha</b> rD,D(rA) <b>e_lhau</b> rD,D8(rA)	Load Halfword Algebraic Load Halfword Algebraic with Update	<a href="#">Section on page 758</a>
<b>e_lhz</b> rD,D(rA) <b>e_lhzu</b> rD,D8(rA)	Load Halfword and Zero Load Halfword and Zero with Update	<a href="#">Section on page 759</a>
<b>e_li</b> rD,LI20	Load Immediate	<a href="#">Section on page 760</a>

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>e_lis</b> rD,UI	Load Immediate Shifted	<a href="#">Section on page 760</a>
<b>e_lmw</b> rD,D8(rA)	Load Multiple Word	<a href="#">Section on page 761</a>
<b>e_lwz</b> rD,D(rA) <b>e_lwzu</b> rD,D8(rA)	Load Word and Zero Load Word and Zero with Update	<a href="#">Section on page 761</a>
<b>e_mcrf</b> crD,crS	Move Condition Register Field	<a href="#">Section on page 763</a>
<b>e_mulli</b> rD,rA,SCI8 <b>e_mull2i</b> rD,SI	Multiply Low Immediate	<a href="#">Section on page 765</a>
<b>e_or2is</b> rD,UI	OR Immediate Shifted	<a href="#">Section on page 765</a>
<b>e_ori[.]</b> rA,rS,SCI8 <b>e_or2i</b> rD,UI	OR Immediate	<a href="#">Section on page 768</a>
<b>e_rlw</b> rA,rS,rB	Rotate Left Word	<a href="#">Section on page 768</a>
<b>e_rlwi</b> rA,rS,SH	Rotate Left Word Immediate	<a href="#">Section on page 768</a>
<b>e_rlwimi</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then Mask Insert	<a href="#">Section on page 769</a>
<b>e_rlwinm</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then AND with Mask	<a href="#">Section on page 770</a>
<b>e_slwi</b> rA,rS,SH	Shift Left Word Immediate	<a href="#">Section on page 761</a>
<b>e_srwi</b> rA,rS,SH	Shift Right Word Immediate	Book E
<b>e_stb</b> rS,D(rA) <b>e_stbu</b> rS,D8(rA)	Store Byte Store Byte with Update	<a href="#">Section on page 773</a>
<b>e_sth</b> rS,D(rA) <b>e_sthu</b> rS,D8(rA)	Store Halfword Store Halfword with Update	<a href="#">Section on page 773</a>
<b>e_stmw</b> rS,D8(rA)	Store Multiple Word	<a href="#">Section on page 774</a>
<b>e_stw</b> rS,D(rA) <b>e_stwu</b> rS,D8(rA)	Store Word Store Word with Update	<a href="#">Section on page 775</a>
<b>e_subfic</b> rD,rA,SCI8 <b>e_subfic.</b> rD,rA,SCI8	Subtract From Immediate Carrying	<a href="#">Section on page 776</a>
<b>e_xori[.]</b> rA,rS,SCI8	XOR Immediate	<a href="#">Section on page 765</a>
<b>icbi</b> rA,rB	Instruction Cache Block Invalidate	Book E
<b>icbt</b> CT,rA,rB	Instruction Cache Block Touch	Book E
<b>isel</b> rD,rA,rB,crb	Integer Select	EIS

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	Book E
<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	Book E
<b>lhbrx</b> rD,rA,rB	Load Halfword Byte-Reverse Indexed	Book E
<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	Book E
<b>lwarx</b> rD,rA,rB	Load Word And Reserve Indexed	Book E
<b>lwbrx</b> rD,rA,rB	Load Word Byte-Reverse Indexed	Book E
<b>lwzx</b> rD,rA,rB <b>lwzux</b> rD,rA,rB	Load Word and Zero Indexed Load Word and Zero with Update Indexed	Book E
<b>mbar</b>	Memory Barrier	Book E
<b>mcrxr</b> crD	Move to Condition Register from Integer Exception Register	Book E
<b>mfc</b> rD	Move From condition register	Book E
<b>mfdc</b> rD,DCRN	Move From Device Control Register	Book E
<b>mfmsr</b> rD	Move From Machine State Register	Book E
<b>mfspir</b> rD,SPRN	Move From Special Purpose Register	Book E
<b>msync</b>	Memory Synchronize	Book E
<b>mtcrf</b> FXM,rS	Move to Condition Register Fields	Book E
<b>mtdc</b> DCRN,rS	Move To Device Control Register	Book E
<b>mtmsr</b> rS	Move To Machine State Register	Book E
<b>mtspir</b> SPRN,rS	Move To Special Purpose Register	Book E
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word	Book E
<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Multiply High Word Unsigned	Book E
<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Multiply Low Word	Book E
<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	NAND	Book E
<b>neg</b> rD,rA <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Negate	Book E



Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	NOR	Book E
<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB	OR	Book E
<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	OR with Complement	Book E
<b>se_add</b> rX,rY	Add	<a href="#">Section on page 741</a>
<b>se_addi</b> rX,OIMM	Add Immediate	<a href="#">Section on page 741</a>
<b>se_andc</b> rX,rY	AND with Complement	<a href="#">Section on page 743</a>
<b>se_andi</b> rX,UI5	AND Immediate	<a href="#">Section on page 743</a>
<b>se_and[.]</b> rX,rY	AND	<a href="#">Section on page 743</a>
<b>se_bc</b> BO16,BI16,BD8	Branch Conditional	<a href="#">Section on page 745</a>
<b>se_bclri</b> rX,UI5	Bit Clear	<a href="#">Section on page 746</a>
<b>se_bctr</b> <b>se_bctrl</b>	Branch to Count Register Branch to Count Register & Link	<a href="#">Section on page 746</a>
<b>se_bgeni</b> rX,UI5	Bit Generate	<a href="#">Section on page 746</a>
<b>se_blr</b> <b>se_blrl</b>	Branch to Link Register Branch to Link Register & Link	<a href="#">Section on page 747</a>
<b>se_bmski</b> rX,UI5	Bit Mask Generate	<a href="#">Section on page 747</a>
<b>se_bseti</b> rX,UI5	Bit Set	<a href="#">Section on page 748</a>
<b>se_b</b> BD8 <b>se_bl</b> BD8	Branch Branch & Link	<a href="#">Section on page 748</a>
<b>se_btsti</b> rX,UI5	Bit Test Immediate	<a href="#">Section on page 744</a>
<b>se_cmp</b> rX,rY	Compare	<a href="#">Section on page 749</a>
<b>se_cmph</b> rX,rY	Compare Halfword	<a href="#">Section on page 750</a>
<b>se_cmphl</b> rX,rY	Compare Halfword Logical	<a href="#">Section on page 751</a>

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>se_cmpi</b> rX,UI5	Compare Immediate	<a href="#">Section on page 749</a>
<b>se_cmpl</b> rX,rY	Compare Logical	<a href="#">Section on page 752</a>
<b>se_cmpli</b> rX,UI5	Compare Logical Immediate	<a href="#">Section on page 752</a>
<b>se_extsb</b> rX	Extend Sign Byte	<a href="#">Section on page 755</a>
<b>se_extsh</b> rX	Extend Sign Halfword	<a href="#">Section on page 755</a>
<b>se_extzb</b> rX	Extend with Zeros Byte	<a href="#">Section on page 756</a>
<b>se_extzh</b> rX	Extend with Zeros Halfword	<a href="#">Section on page 756</a>
<b>se_isync</b>	Instruction Synchronize	<a href="#">Section on page 757</a>
<b>se_lbz</b> rZ,SD4(rX)	Load Byte and Zero (16-bit form)	<a href="#">Section on page 757</a>
<b>se_lhz</b> rZ,SD4(rX)	Load Halfword and Zero (16-bit form)	<a href="#">Section on page 759</a>
<b>se_li</b> rX,UI7	Load Immediate	<a href="#">Section on page 760</a>
<b>se_lwz</b> rZ,SD4(rX)	Load Word and Zero (16-bit form)	<a href="#">Section on page 761</a>
<b>se_mfar</b> rX,arY	Move from Alternate Register	<a href="#">Section on page 762</a>
<b>se_mfctr</b> rX	Move From Count Register	<a href="#">Section on page 762</a>
<b>se_mflr</b> rX	Move From Link Register	<a href="#">Section on page 762</a>
<b>se_mr</b> rX,rY	Move Register	<a href="#">Section on page 762</a>
<b>se_mtar</b> arX,rY	Move to Alternate Register	<a href="#">Section on page 763</a>
<b>se_mtctr</b> rX	Move To Count Register	<a href="#">Section on page 763</a>
<b>se_mtlr</b> rX	Move To Link Register	<a href="#">Section on page 763</a>
<b>se_mullw</b> rX,rY	Multiply Low Word	<a href="#">Section on page 764</a>
<b>se_neg</b> rX	Negate	<a href="#">Section on page 764</a>

Table 146. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>se_not</b> rX	NOT	<a href="#">Section on page 764</a>
<b>se_or</b> rX,rY	OR	<a href="#">Section on page 765</a>
<b>se_slw</b> rX,rY	Shift Left Word	<a href="#">Section on page 770</a>
<b>se_slwi</b> rX,UI5	Shift Left Word Immediate	<a href="#">Section on page 770</a>
<b>se_sraw</b> rX,rY	Shift Right Algebraic Word	<a href="#">Section on page 771</a>
<b>se_srawi</b> rX,UI5	Shift Right Algebraic Word Immediate	<a href="#">Section on page 771</a>
<b>se_srw</b> rX,rY	Shift Right Word	<a href="#">Section on page 772</a>
<b>se_srwi</b> rX,UI5	Shift Right Word Immediate	<a href="#">Section on page 772</a>
<b>se_stb</b> rZ,SD4(rX)	Store Byte (16-bit form)	<a href="#">Section on page 773</a>
<b>se_sth</b> rZ,SD4(rX)	Store Halfword (16-bit form)	<a href="#">Section on page 773</a>
<b>se_stw</b> rZ,SD4(rX)	Store Word (16-bit form)	<a href="#">Section on page 775</a>
<b>se_sub</b> rX,rY	Subtract	<a href="#">Section on page 775</a>
<b>se_subf</b> rX,rY	Subtract From	<a href="#">Section on page 776</a>
<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	Subtract Immediate	<a href="#">Section on page 776</a>
<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB	Shift Left Word	Book E
<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB	Shift Right Algebraic Word	Book E
<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH	Shift Right Algebraic Word Immediate	Book E
<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB	Shift Right Word	Book E
<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Store Byte Indexed Store Byte with Update Indexed	Book E
<b>sthrx</b> rS,rA,rB	Store Halfword Byte-Reverse Indexed	Book E
<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Store Halfword Indexed Store Halfword with Update Indexed	Book E

**Table 146. Instructions listed by mnemonic (continued)**

Mnemonic	Instruction	Reference
<b>stwbrx</b> rS,rA,rB	Store Word Byte-Reverse Indexed	Book E
<b>stwcx.</b> rS,rA,rB	Store Word Conditional Indexed	Book E
<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Store Word Indexed Store Word with Update Indexed	Book E
<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Subtract From	Book E
<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Subtract From Carrying	Book E
<b>tlbivax</b> rA,rB	TLB Invalidate Virtual Address Indexed	Book E
<b>tlbre</b>	TLB Read Entry	Book E
<b>tlbsx</b> rA,rB	TLB Search Indexed	Book E
<b>tlbsync</b>	TLB Synchronize	Book E
<b>tlbwe</b>	TLB Write Entry	Book E
<b>tw</b> TO,rA,rB	Trap Word	Book E
<b>wrtee</b> rA	Write MSR External Enable	Book E
<b>wrteei</b> E	Write MSR External Enable Immediate	Book E
<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	XOR	Book E

## 4.7 Instruction listing

*Table 147* lists instructions defined in Book E, in the PowerPC architecture, and by the EIS. A check mark (√) or text in a column indicates that the instruction is defined or implemented. The EIS-specific instructions are indicated by the name of the APU or architectural extension that defines the instruction.

**Table 147. List of instructions**

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
<b>addc</b> [o][.]	√	√		<b>e_cmpli</b>			VLE
<b>adde</b> [o][.]	√	√		<b>e_crand</b>			VLE
<b>addi</b>	√	√		<b>e_crandc</b>			VLE
<b>addic</b> [.]	√	√		<b>e_creqv</b>			VLE
<b>addis</b>	√	√		<b>e_crnand</b>			VLE
<b>addme</b> [o][.]	√	√		<b>e_crnor</b>			VLE

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
addze[o][.]	√	√		e_cror			VLE
add[o].]	√	√		e_crorc			VLE
andc[.]	√	√		e_crxor			VLE
andi.	√	√		e_lbz			VLE
andis.	√	√		e_lbzu			VLE
and[.]	√	√		e_lha			VLE
b	√	√		e_lhau			VLE
ba	√	√		e_lhz			VLE
bc	√	√		e_lhzu			VLE
bca	√	√		e_li			VLE
bcctr	√	√		e_lis			VLE
bcctrl	√	√		e_lmw			VLE
bcl	√	√		e_lwz			VLE
bcla	√	√		e_lwzu			VLE
bclr	√	√		e_mcrf			VLE
bclrl	√	√		e_mull2i			VLE
bl	√	√		e_mulli			VLE
bla	√	√		e_or2i			VLE
brinc			SPE APU	e_or2is			VLE
cmp	√	√		e_ori[.]			VLE
cmpi	√	√		e_rlw			VLE
cmpl	√	√		e_rlwi			VLE
cmpli	√	√		e_rlwimi			VLE
cntlzw[.]	√	√		e_rlwinm			VLE
crand	√	√		e_slwi			VLE
crandc	√	√		e_srwi			VLE
creqv	√	√		e_stb			VLE
crnand	√	√		e_stbu			VLE
crnor	√	√		e_sth			VLE
cror	√	√		e_sthu			VLE
crorc	√	√		e_stmw			VLE
crxor	√	√		e_stw			VLE
dcba	√	√		e_stwu			VLE
dcbf	√	√		e_subfic			VLE

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
<b>dcbi</b>	√	√		<b>e_subfic.</b>			VLE
<b>dcblc</b>			Cache locking	<b>e_xori[.]</b>			VLE
<b>dcbst</b>	√	√		<b>fabs[.]</b>	√	√	
<b>dcbt</b>	√	√		<b>fadds[.]</b>	√	√	
<b>dcbtls</b>			Cache locking	<b>fadd[.]</b>	√	√	
<b>dcbtst</b>	√	√		<b>fcfid[.]</b>	√	√	
<b>dcbtstls</b>			Cache locking	<b>fcmpo</b>	√	√	
<b>dcbz</b>	√	√		<b>fcmpu</b>	√	√	
<b>divwu[o][.]</b>	√	√		<b>fctidz[.]</b>	√	√	
<b>divw[o][.]</b>	√	√		<b>fctid[.]</b>	√	√	
<b>eciwx</b>		√		<b>fctiwz[.]</b>	√	√	
<b>ecowx</b>		√		<b>fctiw[.]</b>	√	√	
<b>efsabs</b>			Scalar SPFP	<b>fdivs[.]</b>	√	√	
<b>efsadd</b>			Scalar SPFP	<b>fdiv[.]</b>	√	√	
<b>efscfsf</b>			Scalar SPFP	<b>fmadds[.]</b>	√	√	
<b>efscfsi</b>			Scalar SPFP	<b>fmadd[.]</b>	√	√	
<b>efscfuf</b>			Scalar SPFP	<b>fmr[.]</b>	√	√	
<b>efscfui</b>			Scalar SPFP	<b>fmsubs[.]</b>	√	√	
<b>efscmpeq</b>			Scalar SPFP	<b>fmsub[.]</b>	√	√	
<b>efscmpgt</b>			Scalar SPFP	<b>fmuls[.]</b>	√	√	
<b>efscmplt</b>			Scalar SPFP	<b>fmul[.]</b>	√	√	
<b>efscfsf</b>			Scalar SPFP	<b>fnabs[.]</b>	√	√	
<b>efscfsi</b>			Scalar SPFP	<b>fneg[.]</b>	√	√	
<b>efscfsiz</b>			Scalar SPFP	<b>fnmadds[.]</b>	√	√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
efsctuf			Scalar SPFP	fnmadd[.]	√	√	
efsctui			Scalar SPFP	fnmsubs[.]	√	√	
efsctuiz			Scalar SPFP	fnmsub[.]	√	√	
efsddiv			Scalar SPFP	fres[.]	√	√	
efsmul			Scalar SPFP	frsp[.]	√	√	
efsnabs			Scalar SPFP	frsqrte[.]	√	√	
efsneg			Scalar SPFP	fsel[.]	√	√	
efssub			Scalar SPFP	fsqrts[.]	√	√	
efststeg			Scalar SPFP	fsqrt[.]	√	√	
efststgt			Scalar SPFP	fsubs[.]	√	√	
efststlt			Scalar SPFP	fsub[.]	√	√	
emieio	Now mbar	√		icbi	√	√	
eqv[.]	√	√		icblc			Cache locking
evabs			SPE APU	icbt	√		
evaddiw			SPE APU	icbtls			Cache locking
evaddsmiaaw			SPE APU	isel			Integer select
evaddssiaaw			SPE APU	isync	√	√	
evaddumiaaw			SPE APU	lbz	√	√	
evaddusiaaw			SPE APU	lbzu	√	√	
evaddw			SPE APU	lbzux	√	√	
evand			SPE APU	lbzx	√	√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evandc			SPE APU	ld		√	
evcmpeq			SPE APU	ldarx		√	
evcmpgts			SPE APU	ldu		√	
evcmpgtu			SPE APU	ldux		√	
evcmplt			SPE APU	ldx		√	
evcmpltu			SPE APU	lfd	√	√	
evcntlsw			SPE APU	lfdu	√	√	
evcntlzw			SPE APU	lfdx	√	√	
evdivws			SPE APU	lfdx	√	√	
evdivwu			SPE APU	lfs	√	√	
eveqv			SPE APU	lfsu	√	√	
evextsb			SPE APU	lfsux	√	√	
evextsh			SPE APU	lfsx	√	√	
evfsabs			Vector SPFP	lha	√	√	
evfsadd			Vector SPFP	lhau	√	√	
evfscfsf			Vector SPFP	lhaux	√	√	
evfscfsi			Vector SPFP	lhax	√	√	
evfscfuf			Vector SPFP	lhbrx	√	√	
evfscfui			Vector SPFP	lhz	√	√	
evfscmpeq			Vector SPFP	lhzu	√	√	
evfscmpgt			Vector SPFP	lhzux	√	√	



Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evfscmplt			Vector SPFP	lhzx	√	√	
evfsctsf			Vector SPFP	lmw	√	√	
evfsctsi			Vector SPFP	lswi	√	√	
evfsctsiz			Vector SPFP	lswx	√	√	
evfsctuf			Vector SPFP	lwa		√	
evfsctui			Vector SPFP	lwarx	√	√	
evfsctuiz			Vector SPFP	lwaux		√	
evfsdiv			Vector SPFP	lwax		√	
evfsmul			Vector SPFP	lwbrx	√	√	
evfsnabs			Vector SPFP	lwz	√	√	
evfsneg			Vector SPFP	lwzu	√	√	
evfssub			Vector SPFP	lwzux	√	√	
evfststeg			Vector SPFP	lwzx	√	√	
evfststgt			Vector SPFP	mbar	√		
evfststlt			Vector SPFP	mcrf	√	√	
evlidd			SPE APU	mcrfs	√	√	
evliddx			SPE APU	mcrxr	√	√	
evldh			SPE APU	mfapidi	√		
evldhx			SPE APU	mfcrr	√	√	
evldw			SPE APU	mfdcr	√		
evldwx			SPE APU	mffs[.]	√	√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evlhhesplat			SPE APU	mfmsr	√	√	
evlhhesplatx			SPE APU	mfpmr			Performance monitor
evlhhosplat			SPE APU	mfspr	√	√	
evlhhosplatx			SPE APU	mfsr		√	
evlhhusplat			SPE APU	mfsrin		√	
evlhhusplatx			SPE APU	mftb		√	
evlwhe			SPE APU	msync	√		
evlwhex			SPE APU	mtcrf	√	√	
evlw hos			SPE APU	mtdcr	√		
evlw hosx			SPE APU	mtfsb0[.]	√	√	
evlw hou			SPE APU	mtfsb1[.]	√	√	
evlw houx			SPE APU	mtfsfi[.]	√	√	
evlw hsplat			SPE APU	mtfsf[.]	√	√	
evlw hsplatx			SPE APU	mtmsr	√	√	
evlw wsplat			SPE APU	mtmsrd		64-bit only	
evlw wsplatx			SPE APU	mtpmr			Performance monitor
evmergehi			SPE APU	mtspr	√	√	
evmergehilo			SPE APU	mtsr		√	
evmerge lo			SPE APU	mtsr d		√	
evmerge lohi			SPE APU	mtsr din		√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evmhegsmfaa			SPE APU	mtsrin		√	
evmhegsmfan			SPE APU	mulhd.		√	
evmhegsmiaa			SPE APU	mulhdu.		√	
evmhegsmian			SPE APU	mulhwu[.]	√	√	
evmhegumiaa			SPE APU	mulhw[.]	√	√	
evmhegumian			SPE APU	mulld.		√	
evmhesmf			SPE APU	mulldo.		√	
evmhesmf			SPE APU	mulli	√	√	
evmhesmf			SPE APU	mulw[o][.]	√	√	
evmhesmf			SPE APU	nand[.]	√	√	
evmhesmi			SPE APU	neg[o][.]	√	√	
evmhesmia			SPE APU	nor[.]	√	√	
evmhesmiaaw			SPE APU	orc[.]	√	√	
evmhesmianw			SPE APU	ori	√	√	
evmhessf			SPE APU	oris	√	√	
evmhessfa			SPE APU	or[.]	√	√	
evmhessfaaw			SPE APU	rfci	√		
evmhessfanw			SPE APU	rfi	√	√	
evmhessiaaw			SPE APU	rfid		√	
evmhessianw			SPE APU	rfmci			Machine check
evmheumi			SPE APU	rdcl.		√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evmheumia			SPE APU	rldcr.		√	
evmheumiaaw			SPE APU	rldic.		√	
evmheumianw			SPE APU	rldicl.		√	
evmheusiaaw			SPE APU	rldicr.		√	
evmheusianw			SPE APU	rldimi.		√	
evmhogsmfaa			SPE APU	rlwimi[.]	√	√	
evmhogsmfan			SPE APU	rlwinm[.]	√	√	
evmhogsmiaa			SPE APU	rlwnm[.]	√	√	
evmhogsmian			SPE APU	sc	√	√	
evmhogumiaa			SPE APU	se_add			VLE
evmhogumian			SPE APU	se_addi			VLE
evmhosmf			SPE APU	se_andc			VLE
evmhosmfa			SPE APU	se_andi			VLE
evmhosmfaaw			SPE APU	se_and[.]			VLE
evmhosmfanw			SPE APU	se_b			VLE
evmhosmi			SPE APU	se_bc			VLE
evmhosmia			SPE APU	se_bclri			VLE
evmhosmiaaw			SPE APU	se_bctr			VLE
evmhosmianw			SPE APU	se_bctrl			VLE
evmhossf			SPE APU	se_bgeni			VLE
evmhossfa			SPE APU	se_bl			VLE

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evmhossfaaw			SPE APU	se_blr			VLE
evmhossfanw			SPE APU	se_blrl			VLE
evmhossiaaw			SPE APU	se_bmski			VLE
evmhossianw			SPE APU	se_bseti			VLE
evmhoumi			SPE APU	se_btsti			VLE
evmhoumia			SPE APU	se_cmp			VLE
evmhoumiaaw			SPE APU	se_cmph			VLE
evmhoumianw			SPE APU	se_cmphl			VLE
evmhousiaaw			SPE APU	se_cmpi			VLE
evmhousianw			SPE APU	se_cmpl			VLE
evmra			SPE APU	se_cmpli			VLE
evmwhsmf			SPE APU	se_extsb			VLE
evmwhsmfa			SPE APU	se_extsh			VLE
evmwhsmi			SPE APU	se_extzb			VLE
evmwhsmia			SPE APU	se_extzh			VLE
evmwhssf			SPE APU	se_isync			VLE
evmwhssfafa			SPE APU	se_lbz			VLE
evmwhumi			SPE APU	se_lhz			VLE
evmwhumia			SPE APU	se_li			VLE
evmwlsmiaaw			SPE APU	se_lwz			VLE
evmwlsmianw			SPE APU	se_mfar			VLE

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evmwlssiaaw			SPE APU	se_mfctr			VLE
evmwlssianw			SPE APU	se_mflr			VLE
evmwлумi			SPE APU	se_mr			VLE
evmwлумia			SPE APU	se_mtar			VLE
evmwлумiaaw			SPE APU	se_mtctr			VLE
evmwлумianw			SPE APU	se_mtlr			VLE
evmwлusiaaw			SPE APU	se_mullw			VLE
evmwлusianw			SPE APU	se_neg			VLE
evmwsmf			SPE APU	se_not			VLE
evmwsmfa			SPE APU	se_or			VLE
evmwsmfaa			SPE APU	se_slw			VLE
evmwsmfan			SPE APU	se_slwi			VLE
evmwsmi			SPE APU	se_sraw			VLE
evmwsmia			SPE APU	se_srawi			VLE
evmwsmiaa			SPE APU	se_srw			VLE
evmwsmian			SPE APU	se_srwi			VLE
evmwssf			SPE APU	se_stb			VLE
evmwssfа			SPE APU	se_sth			VLE
evmwssfаa			SPE APU	se_stw			VLE
evmwssfаn			SPE APU	se_sub			VLE
evmwumi			SPE APU	se_subf			VLE

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evmwumia			SPE APU	se_subi			VLE
evmwumiaa			SPE APU	se_subi.			VLE
evmwumian			SPE APU	slbia		√	
evnand			SPE APU	slbie		√	
evneg			SPE APU	sldi		√	
evnor			SPE APU	slw[.]	√	√	
evor			SPE APU	srad.		√	
evorc			SPE APU	sradi.		√	
evrlw			SPE APU	srawi[.]	√	√	
evrlwi			SPE APU	sraw[.]	√	√	
evrndw			SPE APU	srd.		√	
evsel			SPE APU	srw[.]	√	√	
evslw			SPE APU	stb	√	√	
evslwi			SPE APU	stbu	√	√	
evsplatfi			SPE APU	stbux	√	√	
evsplatfi			SPE APU	stbx	√	√	
evsrwis			SPE APU	std		√	
evsrwiu			SPE APU	stdcx.		√	
evsrws			SPE APU	stdu		√	
evsrwu			SPE APU	stdux		√	
evstdd			SPE APU	stdx		√	

Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
evstddx			SPE APU	stfd	√	√	
evstdh			SPE APU	stfdu	√	√	
evstdhx			SPE APU	stfdux	√	√	
evstdw			SPE APU	stfdx	√	√	
evstdwx			SPE APU	stfiwx	√	√	
evstwhe			SPE APU	stfs	√	√	
evstwhex			SPE APU	stfsu	√	√	
evstwho			SPE APU	stfsux	√	√	
evstwhox			SPE APU	stfsx	√	√	
evstwwex			SPE APU	sth	√	√	
evstwwex			SPE APU	sthbrx	√	√	
evstwwo			SPE APU	sthu	√	√	
evstwwox			SPE APU	sthux	√	√	
evsubfsmiaaw			SPE APU	sthx	√	√	
evsubfssiaaw			SPE APU	stmw	√	√	
evsubfumiaaw			SPE APU	stswi	√	√	
evsubfusiaaw			SPE APU	stswx	√	√	
evsubfw			SPE APU	stw	√	√	
evsubifw			SPE APU	stwbrx	√	√	
evxor			SPE APU	stwcx.	√	√	
extsb[.]	√	√		stwu	√	√	
extsh[.]	√	√		stwux	√	√	



Table 147. List of instructions (continued)

Mnemonic	Book E	Classic	EIS	Mnemonic	Book E	Classic	EIS
<b>extsw.</b>		64-bit only		<b>stwx</b>	√	√	
<b>e_add16i</b>			VLE	<b>subfc[o][.]</b>	√	√	
<b>e_add2i.</b>			VLE	<b>subfe[o][.]</b>	√	√	
<b>e_add2is</b>			VLE	<b>subfic</b>	√	√	
<b>e_addi</b>			VLE	<b>subfme[o][.]</b>	√	√	
<b>e_addi.</b>			VLE	<b>subfze[o][.]</b>	√	√	
<b>e_addic</b>			VLE	<b>subf[o][.]</b>	√	√	
<b>e_addic.</b>			VLE	<b>sync</b>	Now msync	√	
<b>e_and2i.</b>			VLE	<b>tlbia</b>		√	
<b>e_and2is.</b>			VLE	<b>tlbie</b>		√	
<b>e_andi[.]</b>			VLE	<b>tlbivax</b>	√		
<b>e_b</b>			VLE	<b>tlbre</b>	√		
<b>e_bc</b>			VLE	<b>tlbsx</b>	√		
<b>e_bcl</b>			VLE	<b>tlbsync</b>	√	√	
<b>e_bl</b>			VLE	<b>tlbwe</b>	√		
<b>e_cmp16i</b>			VLE	<b>tw</b>	√	√	
<b>e_cmph</b>			VLE	<b>twi</b>	√	√	
<b>e_cmph16i</b>			VLE	<b>wrtee</b>	√		
<b>e_cmphl</b>			VLE	<b>wrteei</b>	√		
<b>e_cmphl16i</b>			VLE	<b>xori[.]</b>	√	√	
<b>e_cmpi</b>			VLE	<b>xor[.]</b>	√	√	
<b>e_cmpl16i</b>			VLE				

# 5 Interrupts and exceptions

This chapter provides a general description of the Book E exception and interrupt models as they are implemented on ST processors. It identifies and describes the portions of the interrupt model that are defined by the Book E architecture and by the Book E implementation standards (EIS).

*Note: Terminology*

*The Book E architecture has defined additional resources for interrupt handling. As a result, the terms ‘interrupt’ and ‘exception’ differ somewhat from their use in previous ST documentation, such as the Programming Environments Manual. Use of these terms is now as follows:*

- *An interrupt is the action in which the processor saves its context (typically the machine state register (MSR) and next instruction address) and begins execution at a predetermined interrupt handler address with a modified MSR.*
- *An exception is the event that, if enabled, causes the processor to take an interrupt. Book E describes exceptions as being generated by signals from internal and external peripherals, instructions, the internal timer facility, debug events, or error conditions.*

## 5.1 Overview

Book E defines are two categories of interrupts, noncritical and critical, for which separate resources are provided to save state when the interrupt is taken and to restore state when the interrupt handler returns control to the interrupted program.

Using the model provided by the Book E architecture, the EIS defines additional interrupt types which may be implemented on ST Book E devices. These are described in [Table 148](#).

**Table 148. Interrupt types**

Category	Description	Programming resources
<b>Book E defined</b>		
Noncritical interrupts	First-level interrupts that let the processor change program flow to handle conditions generated by external signals, errors, or unusual conditions arising from program execution or from programmable timer-related events. These interrupts are largely identical to those defined by the OEA.	SRR0/SRR1 SPRs and <b>rfi</b> instruction. Asynchronous noncritical interrupts can be masked by the external interrupt enable bit, MSR[EE].
Critical interrupts	Book E–defined. Critical input, watchdog timer, and debug interrupts. these interrupts can be taken during a noncritical interrupt or during regular program flow. Book E defines the critical input, watchdog timer, debug, and machine check interrupts as critical interrupts. The EIS defines additional resources for machine check and debug interrupts.	Critical save and restore SPRs (CSRR0/CSRR1) and the <b>rfci</b> instruction. Critical input and watchdog timer critical interrupts can be masked by the critical enable bit, MSR[CE]. Debug events can be masked by the debug enable bit MSR[DE].

Table 148. Interrupt types

Category	Description	Programming resources
<b>EIS defined (consult implementation documentation to determine whether these interrupts are implemented)</b>		
Machine check interrupt	The EIS-defined machine check APU provides a separate set of resources for the machine check interrupt, which is similar to the Book E–defined critical interrupt type.	Machine check save and restore SPRs (MCSRR0/MCSRR1) and the <b>rfmci</b> instruction. Can be masked by the machine check enable bit, MSR[ME].
Debug interrupt	The EIS-defined debug APU provides a separate set of resources for the debug interrupt, which is similar to the Book E–defined critical interrupt type.	Debug save and restore SPRs (DSRR0/DSRR1) and the <b>rfdi</b> instruction. Can be masked by the machine check enable bit, MSR[DE]. The debug APU extends the Book E debug register model for more detailed control of debug resources.

All interrupts except EIS-defined interrupts are ordered within the two categories of noncritical and critical, such that only one interrupt of each category is reported, and when an interrupt is processed (taken), no program state is lost. Because save/restore register pairs are serially reusable, care must be taken to preserve program state that may be lost when an unordered interrupt is taken. (See [Section 5.10.](#))

All interrupts except the machine check interrupt are context synchronizing as defined in [Section 4.2.3.6: Context synchronization](#). A machine check interrupt acts like a context-synchronizing operation with respect to subsequent instructions; that is, a machine check interrupt need not satisfy items 1 and 2 of [Section 4.2.3.6: Context synchronization](#), but satisfies items 3 and 4.

## 5.2 Els interrupt definitions

This section gives an overview of additions and modifications to the Book E interrupt model defined by the EIS. Specific details are also provided throughout this chapter. Except for the following, the core complex reports exceptions as specified in Book E:

- The machine check exception differs as follows:
  - It is not processed as a critical interrupt, but uses MCSRR0 and MCSRR1 for saving the return address and the MSR in case the machine check is recoverable.
  - Return From Machine Check Interrupt instruction (**rfmci**) is implemented to support the return to the address saved in MCSRR0.
  - A machine check syndrome register, MCSR, logs the cause of the machine check (instead of ESR).

The core complex reports the machine check exception as described in [Section 5.7.2.](#)

- The following interrupts are defined for use with the embedded floating-point and signal-processing (SPE) APUs:
  - SPE/embedded floating-point unavailable interrupt. IVOR32 (SPR 528) contains the vector offset.

See [Section 5.7.17.1: SPE/embedded floating-point APU unavailable interrupt](#).

- Embedded floating-point data interrupt. IVOR33 (SPR 529) contains the vector offset. See [Section 5.7.17.2: Embedded floating-point data interrupt](#).
- Embedded floating-point round interrupt. IVOR34 (SPR 530) contains the vector offset. See [Section 5.7.17.3: Embedded floating-point round interrupt](#).

The following additional bits are defined to support SPE and SPFP exceptions:

- MSR[38] is defined as the vector available bit (SPE). If this bit is clear and software attempts to execute any of the SPE instructions, the SPE unavailable interrupt is taken. If this bit is set, software can execute any SPE instructions.

*Note:* [SPFP instructions require MSR\[SPE\] to be set. An attempt to execute an SPFP instruction when MSR\[SPE\] is 0 causes an SPE APU unavailable interrupt. Section 4.6.1.2: Embedded vector and scalar floating-point APU instructions, lists affected instructions.](#)

- ESR[SPE], the SPE exception bit, is set when the processor reports an exception related to the execution of SPFP or SPE instructions.
- The debug exception implementation does not support IAC3, IAC4, DAC3, and DAC4 comparisons.
- The core complex supports instruction address compare (IAC1 and IAC2) and data address compare (DAC1 and DAC2) for effective addresses only. Real-address support is not provided.
- Some implementations do not support the Book E-defined floating-point unavailable and auxiliary processor unavailable interrupts.
- Data value compare (DVC) debug exceptions are not supported.
- The interrupt priorities differ from those specified in Book E as described in [Section 5.11](#).
- Alignment exceptions. Vector operations can cause alignment exceptions as described in [Section 5.7.6](#).
- Book E and the machine check APU define sources of externally generated interrupts.

### 5.2.1 Recoverability from interrupts

All interrupts except some machine check interrupts are recoverable. The state of the core complex (return address and MSR contents) is saved when a machine check interrupt is taken. The conditions that cause a machine check may or may not prohibit recovery.

## 5.3 Interrupt registers

[Table 149](#) summarizes registers used for interrupt handling. These registers are described in detail in [Chapter 3: Register model](#).

**Table 149. Interrupt registers defined by the PowerPC architecture**

Register	Description
<b>Book E Interrupt Registers</b>	
Save/restore register 0 (SRR0)	On a noncritical interrupt, SRR0 is set to the current or next instruction address. When <b>rfi</b> is executed, instruction execution continues at the address in SRR0. In general, SRR0 contains the address of the instruction that caused the noncritical interrupt or the address of the instruction to return to after a noncritical interrupt is serviced.
Save/restore register 1 (SRR1)	When a noncritical interrupt is taken, MSR contents are placed into SRR1. When <b>rfi</b> is executed, SRR1 contents are placed into the MSR. SRR1 bits that correspond to reserved MSR bits are also reserved. Note that an MSR bit that is reserved may be altered by <b>rfi</b> .
Critical save/restore register 0 (CSRR0)	When a critical interrupt is taken, CSRR0 is set to the current or next instruction address. When <b>rfdi</b> is executed, instruction execution continues at the address in CSRR0. In general, CSRR0 contains the address of the instruction that caused the critical interrupt, or the address of the instruction to return to after a critical interrupt is serviced.
Critical save/restore register 1 (CSRR1)	When a critical interrupt is taken, MSR contents are placed into CSRR1. When <b>rfdi</b> is executed, CSRR1 contents are placed into the MSR. CSRR1 bits that correspond to reserved MSR bits are also reserved. Note that an MSR bit that is reserved may be altered by <b>rfdi</b> .
Data exception address register (DEAR)	DEAR contains the address referenced by a load, store, or cache management instruction that caused an alignment, data TLB miss, or data storage interrupt.
Interrupt vector prefix register (IVPR)	IVPR[32–47] provides the high-order 48 bits of the address of the interrupt handling routine for each interrupt type. The 16-bit vector offsets are concatenated to the right of IVPR to form the address of the interrupt handling routine. IVPR[48–63] are reserved.
Exception syndrome register (ESR)	<p>Provides a syndrome to differentiate between exceptions that can generate the same interrupt type. When one of these types of interrupts is generated, bits corresponding to the specific exception that generated the interrupt are set and all other ESR bits are cleared. Other interrupt types do not affect the ESR. ESR does not need to be cleared by software. <a href="#">Section 3.9.1.8: Exception syndrome register (ESR)</a>, shows ESR bit definitions.</p> <p>An implementation may define additional ESR bits to identify implementation-specific or architected interrupt types; the EIS defines ESR[ILK] and ESR[SPE].</p> <p><b>Note:</b> System software may need to identify the type of instruction that caused the interrupt and examine the TLB entry and ESR to fully identify the exception or exceptions. For example, because both protection violation and byte-ordering exception conditions may be present, and either causes a data storage interrupt, system software would have to look beyond ESR[BO], such as the state of MSR[PR] in SRR1 and the TLB entry page protection bits, to determine if a protection violation also occurred.</p> <p>The EIS defines ESR[56] as the SPE exception bit (SPE). It is set when the processor reports an exception related to the execution of an SPFP or SPE instruction. Note that the EIS definition of the machine check interrupt uses the machine check syndrome register (MCSR) rather than the ESR.</p>

**Table 149. Interrupt registers defined by the PowerPC architecture (continued)**

Register	Description	
Interrupt vector offset registers (IVORs)	Holds the quad-word index from the base address provided by the IVPR for each interrupt type. IVOR0–IVOR15 are provided for defined interrupt types. SPR numbers corresponding to IVOR16–IVOR31 are reserved. IVOR[32–47,60–63] are reserved. SPR numbers for IVOR32–IVOR63 are allocated for implementation-dependent use. (IVOR32–IVOR34 (SPR 528–530) are used by interrupts defined by the EIS.) IVOR assignments are shown below.	
	Book E–defined interrupts IVOR NumberInterrupt Type IVOR0Critical input IVOR1Machine check IVOR2Data storage IVOR3Instruction storage IVOR4External input IVOR5Alignment IVOR6Program IVOR7Floating-point unavailable IVOR8System call IVOR9Auxiliary processor unavailable IVOR10Decrementer IVOR11Fixed-interval timer interrupt IVOR12Watchdog timer interrupt IVOR13Data TLB error IVOR14Instruction TLB error IVOR15Debug IVOR16–IVOR31Reserved	EIS-defined interrupts (IVOR32–IVOR63) IVOR NumberInterrupt Type IVOR32SPE APU unavailable IVOR33Embedded floating-point data IVOR34Embedded floating-point round IVOR35Performance monitor IVOR36Processor doorbell IVOR37Processor doorbell critical
Machine state register (MSR)	MSR[38] is defined as the vector available bit (SPE). It functions as follows: 0: If software attempts to execute an instruction that tries to access the upper word of a 64-bit GPR, an SPE APU unavailable interrupt is taken. 1: Software can execute any embedded floating-point or SPE instructions.	
<b>EIS-Specific Interrupt Registers</b>		
Machine check save/restore register 0 (MCSRR0)	When a machine check interrupt is taken, MCSRR0 is set to the current or next instruction address. When <b>rfmci</b> is executed, instruction execution continues at the address in MCSRR0. In general, MCSRR0 contains the address of the instruction that caused the machine check interrupt, or the address of the instruction to return to after a machine check interrupt is serviced.	
Machine check save/restore register 1 (MCSRR1)	When a machine check interrupt is taken, MSR contents are placed into MCSRR1. When <b>rfmci</b> is executed, MCSRR1 contents are restored to MSR. MCSRR1 bits that correspond to reserved MSR bits are also reserved. Note that an MSR bit that is reserved may be altered by <b>rfmci</b> .	

Table 149. Interrupt registers defined by the PowerPC architecture (continued)

Register	Description
Machine check syndrome register (MCSR)	<p>When a machine check interrupt is taken, MCSR is updated to differentiate among machine check conditions. MCSR also indicates whether a machine check condition is recoverable. ABIST status is logged in MCSR[48–54]. These read-only bits do not initiate machine check (or any other interrupt). An ABIST bit being set indicates an error being detected in the corresponding module.</p> <p>Processors that do not implement the machine check APU use the Book E–defined ESR for this purpose.</p> <p><a href="#">Section 3.9.1.15: Machine check syndrome register (MCSR)</a>, shows MCSR bit definitions.</p>
Machine check address register (MCAR)	<p>When a machine check interrupt is taken, MCAR is updated with the address of the data associated with the machine check. Note that if a machine check interrupt is caused by a signal, the MCAR contents are not meaningful. See <a href="#">Section 3.9.1.14: Machine check address register (MCAR/MCARU)</a>.</p>

## 5.4 Exceptions

Exceptions are caused directly by instruction execution or by an asynchronous event. In either case, the exception may cause one of several types of interrupts to be invoked.

The following examples are of exceptions caused directly by instruction execution:

- An attempt to execute a reserved-illegal instruction (illegal instruction exception-type program interrupt)
- An attempt by an application program to execute a privileged instruction or to access a privileged SPR (privileged instruction exception-type program interrupt)
- In general, an attempt by an application program to access a nonexistent SPR (unimplemented operation instruction exception-type program interrupt). Note the following behavior defined by the EIS:
  - If MSR[PR] = 1 (user mode), SPR bit 5 = 0 (user-accessible SPR), and the SPR number is invalid, an illegal instruction exception is taken.
  - If MSR[PR] = 0 (supervisor mode) and the SPR number is invalid, an illegal instruction exception is taken.
  - If MSR[PR] = 1, SPR bit 5 = 1, and invalid SPR address (supervisor-only SPR), a privileged instruction exception-type program interrupt is taken.
- Execution of a defined instruction using an invalid form (illegal instruction exception-type program interrupt, unimplemented operation exception-type program interrupt, or privileged instruction exception-type program interrupt).
- An attempt to access a location that is either unavailable (instruction or data TLB error interrupt) or not permitted (instruction or data storage interrupt)
- An attempt to access a location with an effective address alignment not supported by the implementation (alignment interrupt)
- Execution of a System Call (**sc**) instruction (system call interrupt)
- Execution of a **trap** instruction whose trap condition is met (trap interrupt type)
- Execution of a floating-point instruction when floating-point instructions are unavailable (floating-point unavailable interrupt)

- Execution of a floating-point instruction that causes a floating-point enabled exception to exist (enabled exception-type program interrupt)
- Execution of a defined instruction that is not implemented (illegal instruction exception or unimplemented operation exception-type program interrupt)
- Execution of an allocated instruction that is not implemented (illegal instruction exception or unimplemented operation exception-type program interrupt)
- Execution of an allocated instruction when the auxiliary instruction is unavailable (auxiliary unavailable interrupt)
- Execution of an allocated instruction that causes an auxiliary enabled exception (enabled exception-type program interrupt)

Invocation of an interrupt is precise, except that if one of the imprecise modes for invoking a floating-point enabled exception-type program interrupt is in effect, the invocation may be imprecise. When the interrupt is invoked imprecisely, the excepting instruction does not appear to complete before the next instruction starts (because the invocation of the interrupt required to complete execution has not occurred).

## 5.5 Interrupt classes

All interrupts except machine check are categorized by two independent characteristics:

- **Critical/noncritical.** Some interrupt types demand immediate attention even if other interrupt types being processed have not had the opportunity to save the machine state (that is, return address and captured state of the MSR). To enable taking a critical interrupt immediately after a noncritical interrupt is taken (that is, before the machine state is saved), two sets of save/restore register pairs are provided. Critical interrupts use CSRR0/CSRR1, and noncritical interrupts use SRR0/SRR1.
- **Asynchronous/synchronous.** Asynchronous interrupts are caused by events external to instruction execution; synchronous interrupts are caused by instruction execution and are either precise or imprecise. [Table 150](#) describes asynchronous and synchronous interrupts.

**Table 150. Asynchronous and synchronous interrupts**

Class	Description
Asynchronous	Caused by events independent from instruction execution. For asynchronous interrupts, the address reported to the interrupt handling routine is the address of the instruction that would have executed next, had the asynchronous interrupt not occurred.



**Table 150. Asynchronous and synchronous interrupts (continued)**

Class	Description
Synchronous, Precise	<p>Caused directly by instruction execution. Synchronous interrupts are precise or imprecise.</p> <p>These interrupts precisely indicate the address of the instruction causing the exception or, for certain synchronous, precise interrupt types, the address of the immediately following instruction. When the execution or attempted execution of an instruction causes a synchronous, precise interrupt, the following conditions exist at the interrupt point:</p> <p>Whether SRR0 or CSRR0 addresses the instruction causing the exception or the next instruction is determined by the interrupt type and status bits.</p> <p>An interrupt is generated such that all instructions before the instruction causing the exception appear to have completed with respect to the executing processor. However, some accesses associated with these preceding instructions may not have been performed with respect to other processors and mechanisms.</p> <p>The exception-causing instruction may appear not to have begun execution (except for causing the exception), may be partially executed, or may have completed, depending on the interrupt type. See <a href="#">Section 5.9</a>.</p> <p>Architecturally, no instruction beyond the exception-causing instruction executed.</p>
Synchronous, Imprecise	<p>Imprecise interrupts may indicate the address of the instruction causing the exception that generated the interrupt or some instruction after that instruction. When execution or attempted execution of an instruction causes an imprecise interrupt, the following conditions exist at the interrupt point.</p> <p>SRR0 or CSRR0 addresses either the exception-causing instruction or some instruction following the exception-causing instruction that generated the interrupt.</p> <p>An interrupt is generated such that all instructions preceding the instruction addressed by SRR0 or CSRR0 appear to have completed with respect to the executing processor.</p> <p>If context synchronization forces the imprecise interrupt due to an instruction that causes another exception that generates an interrupt (for example, alignment or data storage interrupt), SRR0 addresses the interrupt-forcing instruction, which may have partially executed (see <a href="#">Section 5.9</a>).</p> <p>If execution synchronization forces an imprecise interrupt due to an execution-synchronizing instruction other than <b>msync</b> or <b>isync</b>, SRR0 or CSRR0 addresses the interrupt-forcing instruction, which appears not to have begun execution (except for its forcing the imprecise interrupt). If the interrupt is forced by <b>msync</b> or <b>isync</b>, SRR0 or CSRR0 may address <b>msync</b> or <b>isync</b>, or the following instruction.</p> <p>If context or execution synchronization forces an imprecise interrupt, the instruction addressed by SRR0 or CSRR0 may have partially executed (see <a href="#">Section 5.9</a>). No instruction following the instruction addressed by SRR0 or CSRR0 has executed.</p>

### 5.5.1 Requirements for system reset generation

Book E does not specify a system reset interrupt as was defined in the AIM version of the PowerPC architecture. A system reset is typically initiated in one of the following ways:

- Assertion of a signal that resets the internal state of the core complex
- By writing a 1 to DBCR0[34], if MSR[DE] = 1

## 5.6 Interrupt processing

Associated with each kind of interrupt is an interrupt vector, the address of the initial instruction that is executed when an interrupt occurs.

Interrupt processing consists of saving a small part of the processor's state in certain registers, identifying the cause of the interrupt in another register, and continuing execution at the corresponding interrupt vector location. When an exception exists that causes an interrupt to be generated and it has been determined that the interrupt can be taken, the following steps are performed:

1. SRR0 (for noncritical class interrupts) or CSRR0 (for critical class interrupts) or MCSRR0 for machine check interrupts is loaded with an instruction address that depends on the type of interrupt; see the specific interrupt description for details.
2. The ESR or MCSR is loaded with information specific to the exception type. Note that many interrupt types can only be caused by a single type of exception event, and thus do not need nor use an ESR setting to indicate the cause of the interrupt.
3. SRR1 (for noncritical class interrupts) or CSRR1 (for critical class interrupts) or MCSRR1 for machine check interrupts is loaded with a copy of the MSR contents.
4. New MSR values take effect beginning with the first instruction following the interrupt. The MSR is updated as follows:
  - MSR[SPE,WE,EE,PR,FP,FE0,FE1,IS,DS] are cleared by all interrupts.
  - MSR[CE,DE] are cleared by critical class interrupts and unchanged by noncritical class interrupts.
  - MSR[ME] is cleared by machine check interrupts and unchanged by other interrupts.
  - Other defined MSR bits are unchanged by all interrupts.

MSR fields are described in [Section 3.6.1: Machine state register \(MSR\)](#).

5. Instruction fetching and execution resumes, using the new MSR value, at a location specific to the interrupt type (IVPR[32–47] || IVOR $n$ [48–59] || 0b0000)
 

The IVOR $n$  for the interrupt type is described in [Table 151](#). IVPR and IVOR contents are indeterminate upon reset and must be initialized by system software.

Interrupts do not clear reservations obtained with load and reserve instructions. The operating system should do so at appropriate points, such as at process switch.

At the end of a noncritical interrupt handling routine, executing **rfi** causes the MSR to be restored from the SRR1 contents and instruction execution to resume at the address contained in SRR0. Likewise, **rfdi** and **rfdci** perform the same function at the end of critical and machine check interrupt handling routines respectively, using the critical and machine check save/restore registers.

*Note:* *In general, at process switch, due to possible process interlocks and possible data availability requirements, the operating system needs to consider executing the following:*

**stwcx.**—Clears outstanding reservations to prevent pairing a **lwarx** in the old process with a **stwcx** in the new one

**msync**—Ensures that memory operations of an interrupted process complete with respect to other processors before that process begins executing on another processor

**rfdi**, **rfdci**, **rfdci**, or **isync**—Ensures that instructions in the new process execute in the new context

## 5.7 Interrupt definitions

Table 151 summarizes each interrupt type, the various exception types that may cause that interrupt, the interrupt classification, which ESR bits can be set, which MSR bits can mask the interrupt type, and which IVOR is used to specify the vector address.

**Table 151. Interrupt and exception types**

IVOR	Interrupt Type	Exception Type	Exception Class <sup>(1)</sup>	ESR <sup>(2)</sup>	Mask Bits	Notes	Page
IVOR0	Critical input	Critical input	A, C	—	MSR[CE]	(3)	<a href="#">on page 261</a>
IVOR1	Machine check	Machine check	C	—	MSR[ME]	(4),(5)	<a href="#">on page 262</a>
IVOR2	Data storage (DSI)	Access	SP	[SPE],[ST],[FP,AP]	—	(6)	<a href="#">on page 263</a>
		Load reserve or store conditional to write-through required location (W = 1)	SP	[ST]	—	(6)	
		Cache locking	SP	{DLK <sub>0</sub> ,DLK <sub>1</sub> }[DLK,ILK],[ST]	—	(7)	
		Byte ordering	SP	[ST],[FP,AP],BO	—	—	
IVOR3	Instruction storage (ISI)	Access	SP	—	—	—	<a href="#">on page 265</a>
		Byte ordering	SP	BO	—	—	
IVOR4	External input		A	—	MSR[EE]	(3)	<a href="#">on page 266</a>
IVOR5	Alignment		SP	[ST],[FP,AP],[SPE,AP,ST]	—	—	<a href="#">on page 267</a>
IVOR6	Program	Illegal	SP	PIL	—	—	<a href="#">on page 268</a>
		Privileged	SP	PPR,[AP]	—	—	
		Trap	SP	PTR	—	—	
		Floating-point enabled	<b>SP, SI</b>	FP,[PIE]	MSR[FE0,FE1]	(8),(9)	
		Auxiliary processor enabled	SP	AP	—	(9)	
Unimplemented op	SP	PUO,[FP,AP]	—	(11)			
IVOR7	Floating-point unavailable		SP	—	—	—	<a href="#">on page 270</a>
IVOR8	System call		SP	—	—	—	<a href="#">on page 271</a>
IVOR9	Auxiliary processor unavailable		SP	—	—	—	<a href="#">on page 271</a>

**Table 151. Interrupt and exception types (continued)**

IVOR	Interrupt Type	Exception Type	Exception Class <sup>(1)</sup>	ESR <sup>(2)</sup>	Mask Bits	Notes	Page
IVOR10	Decrementer		A	—	MSR[EE], TCR[DIE]	—	<a href="#">on page 271</a>
IVOR11	Fixed interval timer		A	—	MSR[EE], TCR[FIE]	—	<a href="#">on page 272</a>
IVOR12	Watchdog		A, C	—	MSR[CE], TCR[WIE]	—	<a href="#">on page 273</a>
IVOR13	Data TLB error	Data TLB miss	SP	[SPE],[ST], [FP,AP]	—	—	<a href="#">on page 273</a>
IVOR14	Instruction TLB error	Instruction TLB miss	SP	—	—	—	<a href="#">on page 274</a>
IVOR15	Debug	Trap (synchronous)	A, SP, C	—	MSR[DE], DBCR0[IDM]	—	<a href="#">on page 275</a>
		Instruction address compare (synchronous)	A, SP, C	—	MSR[DE], DBCR0[IDM]	—	
		Data address compare (synchronous)	A, SP, C	—	MSR[DE], DBCR0[IDM]	—	
		Instruction complete	SP, C	—	MSR[DE], DBCR0[IDM]	(10)	
		Branch taken	SP, C	—	MSR[DE], DBCR0[IDM]	(10)	
		Return from interrupt	SP, C	—	MSR[DE], DBCR0[IDM]	—	
		Interrupt taken	SI, C	—	MSR[DE], DBCR0[IDM]	—	
Unconditional debug event	SI, C	—	MSR[DE], DBCR0[IDM]	—			
IVOR32	SPE / Embedded FP APU unavailable	SPE APU unavailable	SP	—	—	(11)	<a href="#">on page 275</a>
IVOR33	Embedded FP data	Embedded FP data exception	SP	—	—	(11)	<a href="#">on page 276</a>
IVOR34	Embedded FP round	Embedded FP round exception	SP	—	—	(11)	<a href="#">on page 276</a>

- A = asynchronous, C = critical, SI = synchronous, imprecise, SP = synchronous, precise
- In general, when an interrupt causes an ESR bit or bits to be set (or cleared) as indicated in the table, it also causes all other ESR bits to be cleared. Special rules may apply for implementation-specific ESR bits  
 Legend: xxx (no brackets) means ESR[xxx] is set.  
 [xxx] means ESR[xxx] could be set.  
 [xxx,yyy] means either ESR[xxx] or ESR[yyy] may be set, but never both.  
 {xxx,yyy} means either ESR[xxx] or ESR[yyy] may be set, or possibly both.
- Although not part of Book E, system interrupt controllers commonly provide independent mask and status bits for critical input and external input interrupt sources.



4. Machine check interrupts are not asynchronous or synchronous. See [Section 5.7.2](#).
5. Machine check status information is commonly provided as part of the system implementation but is not part of Book E.
6. Software must examine the instruction and the subject TLB entry to determine the exact cause of the interrupt.
7. Cache locking and cache locking exceptions are implementation-dependent.
8. The precision of the floating-point enabled exception type is controlled by MSR[FE0,FE1], as described in [Table 161](#). See [Section 5.7.7](#). Also, exception status on the exact cause is available in the FPSCR. (See [Section 3.4.2: Floating-point status and control register \(FPSCR\)](#))  
The precision of the auxiliary processor enabled exception type program interrupt is implementation-dependent.
9. Auxiliary processor exception status is commonly provided as part of the implementation and is not part of Book E.
10. Instruction complete and branch taken debug events are defined only for MSR[DE] = 1 for internal debug mode DBCR0[IDM] = 1. In other words, for internal debug mode with MSR[DE] = 0, instruction complete and branch taken debug events cannot occur, no DBSR status bits are set, and no subsequent imprecise debug interrupt can occur.
11. EIS-defined exception

### 5.7.1 Critical input interrupt

A critical input interrupt occurs when no higher priority exception exists, a critical input exception is presented to the interrupt mechanism, and MSR[CE] = 1. The specific definition of a critical input exception is implementation-dependent but is typically caused by assertion of an asynchronous signal that is part of the system. In addition to MSR[CE], implementations may provide other ways to mask the critical input interrupt.

CSRR0, CSRR1, and MSR are updated as shown in [Table 152](#).

**Table 152. Critical input interrupt register settings**

Register	Setting
CSRR0	Set to the effective address of the next instruction to be executed
CSRR1	Set to the MSR contents at the time of the interrupt
MSR	ME is unchanged. All other MSR bits are cleared.

Instruction execution resumes at address IVPR[32–47] || IVOR0[48–59] || 0b0000.

Critical interrupt input signals are level sensitive; to guarantee that the core complex can take a critical input interrupt, the critical input interrupt signal must be asserted until the interrupt is taken. Otherwise, whether the core complex takes an critical interrupt depends on whether MSR[CE] is set when the critical interrupt signal is asserted.

*Note:* To avoid redundant critical input interrupts, software must take any actions required by the implementation to clear any critical input exception status before reenabling MSR[CE].

### 5.7.2 Machine check interrupt

The EIS defines the machine check APU, which differs from the Book E definition of the machine check interrupt as follows:

- Book E defines machine check interrupts as critical interrupts, but the machine check APU treats them as a distinct interrupt type.
- Machine check is no longer a critical interrupt but uses MCSRR0 and MCSRR1 to save the return address and the MSR in case the machine check is recoverable.
- Return from machine check interrupt instruction (**rfmci**) is implemented to support the return to the address saved in MCSRR0.
- An address related to the machine check may be stored in MCAR, according to [Table 153](#).
- A machine check syndrome register, MCSR, is used to log the cause of the machine check (instead of ESR). The MCSR is described in [Table 153](#).

The following general information applies to both the Book E and EIS definitions. A machine check interrupt occurs when no higher priority exception exists, a machine check exception is presented to the interrupt mechanism, and MSR[ME] = 1. Specific causes of machine check exceptions are implementation-dependent, as are the details of the actions taken on a machine check interrupt.

Machine check interrupts are typically caused by a hardware or memory subsystem failure or by an attempt to access an invalid address. They may be caused indirectly by execution of an instruction, but may not be recognized or reported until long after the processor has executed past the instruction that caused the machine check. As such, machine check interrupts are not thought of as synchronous or asynchronous nor as precise or imprecise.

The following general rules apply:

- No instruction after the one whose address is reported to the machine check interrupt handler in MCSRR0 has begun execution.
- The instruction whose address is reported to the machine check interrupt handler in MCSRR0 and all prior instructions may or may not have completed successfully. All instructions certain to complete appear to have done so within the context existing before the machine check interrupt. No further interrupts (other than possible additional machine check interrupts) occur as a result of those instructions.

If MSR[ME] is cleared, the processor enters checkstop state immediately on detecting the machine check condition.

When a machine check interrupt is taken, registers are updated as shown in [Table 153](#).

**Table 153. Machine check interrupt settings**

Register	Setting
CSRR0 <sup>(1)</sup>	Set to an instruction address. As closely as possible, set to the effective address of an instruction that was executing or about to be executing when the machine check exception occurred.
CSRR1 <sup>(1)</sup>	Set to the MSR contents at the time of the interrupt
MSR	UCLE, SPE, WE, CE, EE, PR, FP, ME, FE0, FE1, DE, IS, DS, PMM, and RI are cleared.
ESR	Implementation-dependent. The EIS uses the MCSR rather than the ESR.

**Table 153. Machine check interrupt settings (continued)**

Register	Setting
<b>Machine Check APU Registers</b>	
MCSRR0	On a best-effort basis, the core complex sets this to an effective address of some instruction that was executing or about to be executing when the machine check condition occurred.
MCSRR1	MSR[37–38,46–55,57–59,61–63] are loaded with equivalent MSR bits. All other bits are reserved.
MCAR/ MCARU	When a machine check interrupt is taken, the machine check address register is updated with the address of the data associated with the machine check. Note that if a machine check interrupt is caused by a signal, the MCAR contents are not meaningful. See <a href="#">Section 3.9.1.14: Machine check address register (MCAR/MCARU)</a> . MCARU is an alias to the upper 32 bits of MCAR.
MCSR	Set according to the machine check condition. See <a href="#">Table 20</a> .

1. These registers are used if the machine check APU is not implemented.

Instruction execution resumes at address IVPR[32–47] || IVOR1[48–59] || 0b0000.

*Note:* If a memory subsystem error causes a machine check interrupt, the subsystem may return incorrect data, which may be placed into registers or on-chip caches.

*For implementations on which a machine check interrupt is caused by referring to an invalid physical address, executing **dcbz** or **dcbz** can cause a delayed machine check interrupt by establishing a data cache block associated with an invalid physical address. A machine check interrupt can occur later if and when an attempt is made to write that block to main memory, for example as the result of executing an instruction that causes a cache miss for which the block is the target for replacement or as the result of executing **dcbst** or **dcbf**.*

### 5.7.3 Data storage interrupt

A data storage interrupt (DSI) occurs when no higher priority exception exists and a data storage exception is presented to the interrupt mechanism. [Table 154](#) describes exception conditions for a data storage interrupt as defined by Book E.

**Table 154. Data storage interrupt exception conditions**

Exception	Cause
Read access control exception	Occurs when either of the following conditions exists: <ul style="list-style-type: none"> <li>– In user mode (MSR[PR] = 1), a load or load-class cache management instruction attempts to access a memory location that is not user-mode read enabled (page access control bit UR = 0).</li> <li>– In supervisor mode (MSR[PR] = 0), a load or load-class cache management instruction attempts to access a location that is not supervisor-mode read enabled (page access control bit SR = 0).</li> </ul>

**Table 154. Data storage interrupt exception conditions (continued)**

Exception	Cause
Write access control exception	Occurs when either of the following conditions exists: <ul style="list-style-type: none"> <li>– In user mode (MSR[PR] = 1), a store or store-class cache management instruction attempts to access a location that is not user-mode write enabled (page access control bit UW = 0).</li> <li>– In supervisor mode (MSR[PR] = 0), a store or store-class cache management instruction attempts to access a location that is not supervisor-mode write enabled (page access control bit SW = 0).</li> </ul>
Byte-ordering exception	The implementation cannot access data in the byte order specified by the page's endian attribute. <b>Note:</b> The byte-ordering exception is provided to assist implementations that cannot support dynamically switching byte ordering between consecutive accesses, the byte order for a class of accesses, or misaligned accesses using a specific byte order. Load/store accesses that cross a page boundary such that endianness changes cause a byte-ordering exception.
Cache locking exception	(EIS) The locked state of one or more cache lines has the potential to be altered. This exception is implementation-dependent. A cache locking exception occurs with the execution of <b>icbtls</b> , <b>icblc</b> , <b>dcbtls</b> , <b>dcbtstls</b> , or <b>dcblc</b> when (MSR[PR] = 1) and (MSR[UCLE] = 0). ESR is set as follows: <ul style="list-style-type: none"> <li>– For <b>icbtls</b> and <b>icblc</b>, ESR[ILK] is set.</li> <li>– For <b>dcbtls</b>, <b>dcbtstls</b>, or <b>dcblc</b>, ESR[DLK] is set. Book E refers to this as a cache-locking exception.</li> </ul>
Storage synchronization exception	Occurs when either of the following conditions exists: <ul style="list-style-type: none"> <li>– An attempt is made to execute a load and reserve or store conditional instruction from or to a location that is write-through required or caching inhibited. (If the interrupt does not occur, the instruction executes correctly.)</li> <li>– A store conditional instruction produces an effective address for which a normal store would cause a data storage interrupt but the processor does not have the reservation from a load and reserve instruction. Book E states that it is implementation-dependent whether a data storage interrupt occurs. The EIS defines that the data storage interrupt is taken.</li> </ul>

Instructions **icbt**, **dcbt**, **dcbtst**, and **dcba**, and **lswx** or **stswx** with a length of zero cannot cause a data storage interrupt, regardless of the effective address.

*Note:* **icbi** and **icbt** are treated as loads from the addressed byte with respect to address translation and protection. They use MSR[DS], not MSR[IS], to determine translation for their operands. Instruction storage interrupts and instruction TLB error interrupts are associated with instruction fetching and not execution. Data storage interrupts and data TLB error interrupts are associated with the execution of instruction cache management instructions.

When a data storage interrupt occurs, the processor suppresses execution of the instruction causing the data storage exception.

SRR0, SRR1, ESR, MSR, and DEAR, are updated as follows:



**Table 155. Data Storage Interrupt Register Settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the interrupt
SRR1	Set to the MSR contents at the time of the interrupt
ESR	FP Set if the instruction causing the interrupt is a floating-point load or store; otherwise cleared ST Set if the instruction causing the interrupt is a store or store-class cache management instruction; otherwise cleared DLK DLK is set when a DSI occurs because <b>dcbtls</b> , <b>dcbtstls</b> , or <b>dcblc</b> is executed in user mode and MSR[UCLE] = 0. AP Set if the instruction causing the interrupt is an auxiliary processor load or store; otherwise cleared BO Set if the instruction caused a byte-ordering exception; otherwise cleared All other defined ESR bits are cleared.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
DEAR	Set to the effective address of a byte that lies both within the range of bytes being accessed by the access or cache management instruction and within the page whose access caused the exception

Instruction execution resumes at address IVPR[32–47] || IVOR2[48–59] || 0b0000.

### 5.7.4 Instruction storage interrupt

An instruction storage interrupt occurs when no higher priority exception exists and an instruction storage exception is presented to the interrupt mechanism. Instruction storage exception conditions are described in [Table 156](#).

**Table 156. Instruction storage interrupt exception conditions**

Exception	Cause
Execute access control exception	In user mode, an instruction fetch attempts to access a memory location that is not user mode execute enabled (page access control bit UX = 0). In supervisor mode, an instruction fetch attempts to access a memory location that is not supervisor mode execute enabled (page access control bit SX = 0).
Byte-ordering exception	The implementation cannot fetch the instruction in the byte order specified by the page’s endian attribute. The EIS defines that accesses that cross a page boundary such that endianness changes cause a byte-ordering exception.

Note that Book E provides this exception to assist implementations that cannot dynamically switch byte ordering between consecutive accesses, do not support the byte order for a class of accesses, or do not support misaligned accesses using a specific byte order.

When an instruction storage interrupt occurs, the processor suppresses execution of the instruction causing the exception.

SRR0, SRR1, MSR, and ESR are updated as shown in [Table 157](#).

**Table 157. Instruction storage interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the instruction storage interrupt
SRR1	Set to the MSR contents at the time of the interrupt
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
ESR	BO is set if the instruction fetch caused a byte-ordering exception; otherwise cleared. All other defined ESR bits are cleared.

*Note:* Permissions violation and byte-ordering exceptions are not mutually exclusive. Even if ESR[BO] is set, system software must examine the TLB entry accessed by the fetch to determine whether a permissions violation also may have occurred.

Instruction execution resumes at address IVPR[32–47] || IVOR3[48–59] || 0b0000.

### 5.7.5 External input interrupt

An external input interrupt occurs when no higher priority exception exists, an external input exception is presented to the interrupt mechanism, and MSR[EE] = 1. The specific definition of an external input exception is implementation-dependent and is typically caused by assertion of an asynchronous signal that is part of the processing system.

To guarantee that the core complex can take an external interrupt, the external interrupt pin must be asserted until the interrupt is taken. Otherwise, whether the external interrupt is taken depends on whether MSR[EE] is set when the external interrupt signal is asserted.

In addition to MSR[EE], implementations may provide other ways to mask this interrupt.

SRR0, SRR1, and MSR are updated as shown in [Table 158](#).

**Table 158. External input interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the next instruction to be executed
SRR1	Set to the MSR contents at the time of the interrupt
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.

Instruction execution resumes at address IVPR[32–47] || IVOR4[48–59] || 0b0000.

*Note:* To avoid redundant external input interrupts, software must take any actions required to clear any external input exception status before reenabling MSR[EE].

### 5.7.6 Alignment interrupt

An alignment interrupt occurs when no higher priority exception exists and an alignment exception is presented to the interrupt mechanism. An alignment exception may occur when an implementation cannot perform a data access for one of the following reasons:

- The operand of a load or store is not aligned.
- The instruction is a move assist, load multiple, or store multiple.
- A **dcbz** operand is in write-through-required or caching-inhibited memory, or **dcbz** is executed in an implementation with no data cache or a write-through data cache.
- The operand of a store, except store conditional, is in write-through required memory.

The EIS defines the following alignment exception conditions:

- Execution of a **dcbz** references a page marked as write-through or cache inhibited.
- A load multiple word instruction (**lmw**) reads an address that is not a multiple of four.
- A **lwarx** or **stwcx**. instruction references an address that is not a multiple of four.
- SPFP and SPE APU instructions are not aligned on a natural boundary. A natural boundary is defined by the size of the data element being accessed.
- A vector operation reports an exception if the physical address of the following instructions is not aligned to the 64-bit boundary: **evldd**, **evlddx**, **evldw**, **evldwx**, **evldh**, **evldhx**, **evstd**, **evstdx**, **evstdw**, **evstdwx**, **evstdh**, and **evstdhx**. [Table 159](#) describes additional ESR settings.

For **lmw** and **stmw** with a non-word-aligned operand and for load and reserve and store conditional instructions with a misaligned operand, an implementation may yield boundedly undefined results instead of causing an alignment interrupt. A store conditional to a write-through required location may either cause an alignment or data storage interrupt or may correctly execute the instruction. For all other cases listed above, an implementation may execute the instruction correctly instead of causing an alignment interrupt. For **dcbz**, correct execution means clearing each byte of the block in main memory.

*Note: Book E does not support use of an misaligned effective address by load and reserve and store conditional instructions. If an misaligned effective address is specified, the alignment interrupt handler should treat the instruction as a programming error and must not attempt to emulate the instruction.*

When an alignment interrupt occurs, the processor suppresses the execution of the instruction causing the alignment exception.

SRR0, SRR1, MSR, DEAR, and ESR are updated as shown in [Table 159](#).

**Table 159. Alignment interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the alignment interrupt
SRR1	Set to the MSR contents at the time of the interrupt
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.

**Table 159. Alignment interrupt register settings (continued)**

Register	Setting
DEAR	Set to the EA of a byte that is both within the range of the bytes being accessed by the memory access or cache management instruction, and within the page whose access caused the alignment exception
ESR	FP Set if the instruction causing the interrupt is a floating-point load or store; otherwise cleared ST Set if the instruction causing the interrupt is a store; otherwise cleared AP Set if the instruction causing the interrupt is an auxiliary processor load or store; otherwise cleared The following bits may be affected for vector alignment exception conditions: SPE Set AP Set (May not be supported on all processors) ST Set only if the instruction causing the exception is a store and is cleared for a load All other defined ESR bits are cleared.

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR5[48-59] \parallel 0b0000$ .

### 5.7.7 Program interrupt

A program interrupt occurs when no higher priority exception exists and a program exception is presented to the interrupt mechanism. A program interrupt is caused when any of the following exceptions occurs during execution of an instruction.

**Table 160. Program interrupt exception conditions**

Exception	Cause
Floating-point enabled exception	Caused when $(MSR[FE0] \mid MSR[FE1]) \& FPSCR[FEX] = 1$ . $FPSCR[FEX]$ is set by the execution of a floating-point instruction that causes an enabled exception, including the case of a Move to FPSCR instruction that causes an exception bit and the corresponding enable bit both to be 1. Note that in this context, the term 'enabled exception' refers to the enabling provided by FPSCR control bits. See <a href="#">Section 3.4.2: Floating-point status and control register (FPSCR)</a> . Whether the interrupt is precise or imprecise is determined by $MSR[FE0,FE1]$ , as described in <a href="#">Table 20</a> .
Auxiliary processor enabled exception	Implementation dependent

**Table 160. Program interrupt exception conditions (continued)**

Exception	Cause
Illegal instruction exception	<p>Always occurs when execution of any of the following kinds of instructions is attempted.</p> <ul style="list-style-type: none"> <li>– A reserved-illegal instruction</li> <li>– In user mode, an <b>mtspr</b> or <b>mfspir</b> that specifies an SPRN value with SPRN[5] = 0 (user-mode accessible) that represents an unimplemented SPR</li> <li>– (EIS) If an invalid SPR address is accessible only in supervisor mode and the processor is in supervisor mode (MSR[PR] = 0), results are undefined.</li> <li>– (EIS) If the invalid SPR address is accessible only in the supervisor mode and the processor is in user mode (MSR[PR] = 1), a privileged instruction exception is taken.</li> </ul> <p>May occur when execution is attempted of any of the following kinds of instructions. If the exception does not occur, the alternative is shown in parentheses. See the user’s manual for the implementation.</p> <ul style="list-style-type: none"> <li>– An instruction that is in invalid form (boundedly undefined results).</li> <li>– An <b>lswx</b> instruction for which <b>rA</b> or <b>rB</b> is in the range of registers to be loaded (boundedly undefined results)</li> <li>– A reserved no-op instruction (no-operation performed is preferred).</li> <li>– A defined or allocated instruction that is not implemented (unimplemented operation exception). Unimplemented Book E instructions take an illegal instruction exception.</li> <li>– The EIS defines that an attempt to execute a 64-bit Book E instruction causes an illegal instruction exception.</li> </ul>
Privileged instruction exception	<p>Occurs when MSR[PR] = 1 and execution is attempted of any of the following:</p> <ul style="list-style-type: none"> <li>– A privileged instruction</li> <li>– An <b>mtspr</b> or <b>mfspir</b> instruction that specifies a privileged SPR (SPRN[5] = 1)</li> <li>– (EIS) An <b>mtpmr</b> or <b>mfpmr</b> instruction that specifies a privileged PMR (PMRN[5] = 1)</li> </ul>
Trap exception	Occurs when any of the conditions specified in a trap instruction are met.
Unimplemented operation exception	May occur when a defined or allocated instruction is encountered that is not implemented. Otherwise an illegal instruction exception occurs. See the reference manual for the implementation.

Whether a floating-point enabled interrupt is precise or imprecise is determined by MSR[FE0,FE1], as described in [Table 161](#).

**Table 161. MSR[FE0,FE1] settings**

FE0,FE1	Description
01,10	Imprecise. When such a program interrupt is taken, if the address saved in SRR0 is not that of the instruction that caused the exception (that is, the instruction that caused FPSCR[FEX] to be set), ESR[PIE] is set. Note that some implementations may ignore these bit settings and treat all affected interrupts as precise.
11	Precise.
00	The interrupt is masked and the interrupt subsequently occurs if and when floating-point enabled exception-type program interrupts are enabled by setting either or both FE0,FE1 and also causes ESR[PIE] to be set.

SRR0, SRR1, MSR, and ESR are updated as shown in [Table 162](#).

**Table 162. Program interrupt register settings**

Register	Description
SRR0	For all program interrupts except an enabled exception when in an imprecise mode (see <a href="#">Table 164</a> ), set to the EA of the instruction that caused the interrupt. For an imprecise enabled exception, set to the EA of the excepting instruction or of some subsequent instruction that has not been executed (in which case ESR[PIE] is set). If the instruction is <b>msync</b> or <b>isync</b> , SRR0 does not point more than 4 bytes beyond the <b>msync</b> or <b>isync</b> . If FPSCR[FEX] = 1 but both MSR[FE0,FE1] = 00, an enabled exception-type program interrupt occurs before or at the next synchronizing event if [FE0,FE1] are altered by any instruction so that the expression (MSR[FE0]   MSR[FE1]) & FPSCR[FEX] is 1. When this occurs, ESR[PIE] is set and SRR0 is loaded with the EA of the instruction that would have executed next, not with the EA of the instruction that modified MSR causing the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
ESR	PIL Set if an illegal instruction exception-type program interrupt; otherwise cleared. PPR Set if a privileged instruction exception-type program interrupt; otherwise cleared. PTR Set if a trap exception-type program interrupt; otherwise cleared. PUO Set if an unimplemented operation exception-type program interrupt; otherwise cleared. FP Set if the instruction causing the interrupt is a floating-point instruction; otherwise cleared. PIE Set if a floating-point enabled exception-type program interrupt, and the address saved in SRR0 is not the address of the instruction causing the exception (that is, the instruction that caused FPSCR[FEX] to be set); otherwise cleared. AP Set if the instruction causing the interrupt is an auxiliary processor instruction; otherwise cleared. All other defined ESR bits are cleared.

Instruction execution resumes at address IVPR[32–47] || IVOR6[48–59] || 0b0000.

### 5.7.8 Floating-point unavailable interrupt

A floating-point unavailable interrupt occurs when no higher priority exception exists, an attempt is made to execute a floating-point instruction (including floating-point loads, stores, and moves), and MSR[FP] = 0.

When a floating-point unavailable interrupt occurs, the processor suppresses execution of the instruction causing the floating-point unavailable interrupt.

SRR0, SRR1, and MSR are updated as shown in [Table 163](#).

**Table 163. Floating-point unavailable interrupt register settings**

Register	Description
SRR0	Set to the effective address of the instruction that caused the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.

Instruction execution resumes at address  $IVPR[32-47] || IVOR7[48-59] || 0b0000$ .

### 5.7.9 System call interrupt

A system call interrupt occurs when no higher priority exception exists and a System Call (**sc**) instruction is executed. SRR0, SRR1, and MSR are updated as shown in [Table 164](#).

**Table 164. System call interrupt register settings**

Register	Description
SRR0	Set to the effective address of the instruction after the <b>sc</b> instruction.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.

Instruction execution resumes at address  $IVPR[32-47] || IVOR8[48-59] || 0b0000$ .

### 5.7.10 Auxiliary processor unavailable interrupt

An auxiliary processor unavailable interrupt occurs when no higher priority exception exists, an attempt is made to execute an auxiliary processor instruction (including auxiliary processor loads, stores, and moves), the target auxiliary processor is present on the implementation, and the auxiliary processor is configured as unavailable. Details of the auxiliary processor and its configuration are implementation-dependent. See the reference manual for the implementation.

When an auxiliary processor unavailable interrupt occurs, the processor suppresses execution of the instruction causing the auxiliary processor unavailable interrupt.

Registers SRR0, SRR1, and MSR are updated as shown in [Table 165](#).

**Table 165. Auxiliary processor unavailable interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction that caused the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.

Instruction execution resumes at address  $IVPR[32-47] || IVOR9[48-59] || 0b0000$ .

### 5.7.11 Decrementer Interrupt

A decrementer interrupt occurs when no higher priority exception exists, a decrementer exception exists ( $TSR[DIS] = 1$ ) & the interrupt is enabled ( $TCR[DIE] = 1$  and  $MSR[EE] = 1$ ).

MSR[EE] also enables external input and fixed-interval timer interrupts.

SRR0, SRR1, MSR, and TSR are updated as shown in [Table 166](#).

**Table 166. Decrementer interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the next instruction to be executed.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
TSR	DIS is set.

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR10[48-59] \parallel 0b0000$ .

*Note:* To avoid redundant decrementer interrupts, before reenabling MSR[EE], the interrupt handling routine must clear TSR[DIS] by writing a word to TSR using **mtspr** with a 1 in any bit position to be cleared and 0 in all others. The data written to the TSR is not direct data, but a mask. Writing a 1 to this bit causes it to be cleared; writing a 0 has no effect.

### 5.7.12 Fixed-interval timer interrupt

A fixed-interval timer interrupt occurs when no higher priority exception exists, a fixed-interval timer exception exists (TSR[FIS] = 1), and the interrupt is enabled (TCR[FIE] = 1 and MSR[EE] = 1).

The fixed-interval timer period is determined by TCR[FP], which, when concatenated with TCR[FPEXT], specifies one of 64 bit locations of the time base used to signal a fixed-interval timer exception on a transition from 0 to 1.

TCR[FPEXT],TCR[FP] = 000000 selects TBU[32]. TCR[FPEXT],TCR[FP] = 111111 selects TBL[63].

*Note:* MSR[EE] also enables external input and decrementer interrupts.

SRR0, SRR1, MSR, and TSR are updated as shown in [Table 167](#).

**Table 167. Fixed-interval timer interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the next instruction to be executed.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
TSR	FIS is set.

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR11[48-59] \parallel 0b0000$ .

*Note:* To avoid redundant fixed-interval timer interrupts, before reenabling MSR[EE], the interrupt handling routine must clear TSR[FIS] by writing a word to TSR using **mtspr** with a 1 in any bit position to be cleared and 0 in all others. The data written to the TSR is not direct data, but a mask. Writing a 1 causes the bit to be cleared; writing a 0 has no effect.



### 5.7.13 Watchdog timer interrupt

A watchdog timer interrupt occurs when no higher priority exception exists, a watchdog timer exception exists (TSR[WIS] = 1), and the interrupt is enabled (TCR[WIE] = 1 and MSR[CE] = 1).

MSR[CE] also enables the critical input interrupt.

CSRR0, CSRR1, MSR, and TSR are updated as shown in [Table 168](#).

**Table 168. Watchdog timer interrupt register settings**

Register	Setting
CSRR0	Set to the effective address of the next instruction to be executed.
CSRR1	Set to the MSR contents at the time of the interrupt.
MSR	ME is unchanged; all other MSR bits are cleared.
TSR	WIS is set.

Instruction execution resumes at address IVPR[32–47] || IVOR12[48–59] || 0b0000.

*Note:* To avoid redundant watchdog timer interrupts, before reenabling MSR[CE], the interrupt handling routine must clear TSR[WIS] by writing a word to TSR using *mtspr* with a 1 in any bit position to be cleared and 0 in all others. The data written to the TSR is not direct data, but a mask. Writing a 1 to this bit causes it to be cleared; writing a 0 has no effect.

### 5.7.14 Data tlb error interrupt

A data TLB error interrupt occurs when no higher priority exception exists and the exception described in [Table 169](#) is presented to the interrupt mechanism.

**Table 169. Data tlb error interrupt exception conditions**

Exception	Description
Data TLB miss exception	Virtual addresses associated with an instruction fetch do not match any valid TLB entry.

If a store conditional instruction produces an effective address for which a normal store would cause a data TLB error interrupt, but the processor does not have the reservation from a load and reserve instruction, Book E defines it as implementation-dependent whether a data TLB error interrupt occurs. The EIS defines that the interrupt is taken.

When a data TLB error interrupt occurs, the processor suppresses execution of the instruction causing the data TLB error exception.

SRR0, SRR1, MSR, DEAR, and ESR are updated as shown in [Table 170](#).

**Table 170. Data tlb error interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the data TLB error interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
DEAR	Set to the EA of a byte that is both within the range of the bytes being accessed by the memory access or cache management instruction and within the page whose access caused the data TLB error exception.
ESR	STSet if the instruction causing the interrupt is a store, <b>dcbi</b> , or <b>dcbz</b> instruction; otherwise cleared. FPSet if the instruction causing the interrupt is a floating-point load or store; otherwise cleared. APSet if the instruction causing the interrupt is an auxiliary processor load or store; otherwise cleared. All other defined ESR bits are cleared.
MAS <sub>n</sub>	See <a href="#">Table 193</a> .

[Table 192](#) shows MAS register settings for data and instruction TLB error interrupts. [Section 6.4.11.3: MAS register updates for exceptions, tlbsx, and tlbre](#), describes how these values are set as defined by the EIS.

Instruction execution resumes at address IVPR[32–47] || IVOR13[48–59] || 0b0000.

### 5.7.15 Instruction tlb error interrupt

An instruction TLB error interrupt occurs when no higher priority exception exists and the exception described in [Table 171](#) is presented to the interrupt mechanism.

**Table 171. Instruction TLB error interrupt exception conditions**

Exception	Description
Instruction TLB miss exception	The virtual addresses associated with a fetch do not match any valid TLB entry.

When an instruction TLB error interrupt occurs, the processor suppresses execution of the instruction causing the instruction TLB miss exception.

SRR0, SRR1, and MSR are updated as shown in [Table 172](#).

**Table 172. Instruction TLB error interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the instruction TLB error interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
MAS <sub>n</sub>	See <a href="#">Table 192</a> .

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR14[48-59] \parallel 0b0000$ .

### 5.7.16 Debug interrupt

A debug interrupt occurs when no higher priority interrupt exists, a debug exception exists in the DBSR, and debug interrupts are enabled ( $DBCRO[IDM] = 1$  and  $MSR[DE] = 1$ ). A debug exception occurs when a debug event causes a corresponding DBSR bit to be set.

**Table 173. Debug interrupt register settings**

Register	Setting
CSRR0	<p>For debug exceptions that occur while debug interrupts are enabled (<math>DBCRO[IDM] = 1</math> and <math>MSR[DE] = 1</math>), CSRR0 is set as follows:</p> <ul style="list-style-type: none"> <li>– For instruction address compare (IAC registers), data address compare (DAC1R, DAC1W, DAC2R, and DAC2W), data value compare (DVC1 and DVC2), trap (TRAP), or branch taken (BRT) debug exceptions, set to the address of the instruction causing the debug interrupt.</li> <li>– For instruction complete (ICMP) debug exceptions, set to the address of the instruction that would have executed after the one that caused the debug interrupt.</li> <li>– For unconditional debug event (UDE) debug exceptions, set to the address of the instruction that would have executed next if the debug interrupt had not occurred.</li> <li>– For interrupt taken (IRPT) debug exceptions, set to the interrupt vector value of the interrupt that caused the interrupt taken debug event.</li> <li>– For return from interrupt (RET) debug exceptions, set to the address of the instruction that would have executed after the <b>rfi</b>, <b>rfdci</b>, or <b>rfmci</b> that caused the debug interrupt.</li> <li>– For debug exceptions that occur while debug interrupts are disabled (<math>DBCRO[IDM] = 0</math> or <math>MSR[DE] = 0</math>), a debug interrupt occurs at the next synchronizing event if <math>DBCRO[IDM]</math> and <math>MSR[DE]</math> are modified such that they are both set and if the debug exception status is still set in the DBSR. When this occurs, CSRR0 holds the address of the instruction that would have executed next, not the address of the instruction that modified <math>DBCRO</math> or <math>MSR</math> and thus caused the interrupt.</li> </ul>
CSRR1	Set to the MSR contents at the time of the interrupt.
MSR	ME is unchanged. All other MSR bits are cleared.
DBSR	Set to indicate type of debug event. (See <a href="#">Section 3.13.2</a> .)

CSRR0, CSRR1, MSR, and DBSR are updated as shown in [Table 173](#).

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR15[48-59] \parallel 0b0000$ .

### 5.7.17 EIS-defined interrupts

The interrupts in this section are defined by the EIS.

#### 5.7.17.1 SPE/embedded floating-point APU unavailable interrupt

An SPE APU unavailable interrupt is taken if  $MSR[SPE]$  is cleared and an SPE, embedded scalar double-precision or embedded vector single-precision floating-point instruction is executed. It is not used by the embedded scalar single-precision floating-point APU.

When an SPE unavailable interrupt occurs, the processor suppresses execution of the instruction causing the interrupt. [Table 174](#) describes register settings. If the SPE/embedded floating-point unavailable interrupt occurs, the processor suppresses

execution of the instruction causing the exception. The SRR0, SRR1, MSR, and ESR registers are modified as shown in [Table 174](#).

**Table 174. SPE/embedded floating-point APU unavailable interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other bits are cleared.
ESR	SPE (bit 24) is set. All other ESR bits are cleared.

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

### 5.7.17.2 Embedded floating-point data interrupt

An embedded floating-point data interrupt is generated in the following cases:

- SPEFSCR[FINVE] = 1 and either SPEFSCR[FINVH,FINV] = 1
- SPEFSCR[FDBZE] = 1 and either SPEFSCR[FDBZH,FDBZ] = 1
- SPEFSCR[FUNFE] = 1 and either SPEFSCR[FUNFH,FUNF] = 1
- SPEFSCR[FOVFE] = 1 and either SPEFSCR[FOVFH,FOVF] = 1

Note that although SPEFSCR status bits can be updated by using **mtspr**, interrupts occur only if they are set as the result of an arithmetic operation.

When an embedded floating-point data interrupt occurs, the processor suppresses execution of the instruction causing the interrupt. [Table 175](#) shows register settings.

**Table 175. Embedded floating-point data interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction causing the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other bits are cleared.
ESR	SPE (bit 24) is set. All other ESR bits are cleared.
SPEFSCR	One or more of the FINVH, FINV, FDBZH, FDBZ, FUNFH, FUNF, FOVFH, or FOVF bits are set to indicate the interrupt type.

Instruction execution resumes at address IVPR[32–47] || IVOR33[48–59] || 0b0000.

### 5.7.17.3 Embedded floating-point round interrupt

The embedded floating-point round interrupt is taken on any of the following conditions:

- SPEFSCR[FINXE] = 1 and any of the SPEFSCR[FGH,FXH,FG,FX] bits = 1
- SPEFSCR[FRMC] = 0b10 (+∞)
- SPEFSCR[FRMC] = 0b11 (−∞)

Note that although these SPEFSCR status bits can be updated by using an **mtspr**[SPEFSCR], interrupts occur only if they are set as the result of an arithmetic operation.

When an embedded floating-point round interrupt occurs, the unrounded (truncated) result is placed in the target register. [Table 176](#) describes register settings.

**Table 176. Embedded floating-point round interrupt register settings**

Register	Setting
SRR0	Set to the effective address of the instruction following the instruction causing the interrupt.
SRR1	Set to the MSR contents at the time of the interrupt.
MSR	CE, ME, and DE are unchanged. All other MSR bits are cleared.
ESR	SPE (bit 24) is set. All other ESR bits are cleared.
SPEFSCR	FGH, FXH, FG, FX, and FRMC are set appropriately to indicate the interrupt type.

Instruction execution resumes at address  $IVPR[32-47] \parallel IVOR34[48-59] \parallel 0b0000$ .

## 5.8 Performance monitor interrupt

The performance monitor provides a performance monitor interrupt that is triggered by an enabled condition or event. An enabled condition or event is as follows:

A PMC<sub>x</sub> register overflow condition occurs with the following settings:

- PMLC<sub>x</sub>[CE] = 1; that is, for the given counter the overflow condition is enabled.
- PMC<sub>x</sub>[32] = 1; that is, the given counter indicates an overflow.

For a performance monitor interrupt to be signaled on an enabled condition or event, PMGC0[PMIE] must be set.

The performance monitor can also freeze the performance monitor counters triggered by an enabled condition or event. For the performance monitor counters to freeze on an enabled condition or event, PMGC0[FCECE] must be set.

Although the interrupt condition could occur with MSR[EE] = 0, the interrupt cannot be taken until MSR[EE] = 1. If a counter overflows while PMGC0[FCECE] = 0, PMLC<sub>a</sub>[CE] = 1, and MSR[EE] = 0, it is possible for the counter to wrap around to all zeros again without the performance monitor interrupt being taken.

The priority of the performance monitor interrupt is below that of the fixed-interval interrupt and above that of the decremter interrupt.

## 5.9 Partially executed instructions

In general, the PowerPC architecture permits load and store instructions to be partially executed, interrupted, and then restarted from the beginning upon return from the interrupt. To guarantee that a particular load or store instruction completes without being interrupted and restarted, software must mark the memory as guarded and use an elementary (non-string or non-multiple) load or store aligned on an operand-sized boundary.

To guarantee that load and store instructions can, in general, be restarted and completed correctly without software intervention, the following rules apply when an execution is partially executed and then interrupted:

- For an elementary load, no part of a target register **rD** or **frD** has been altered.
- For update forms of load or store, the update register, **rA**, will not have been altered.

The following effects are permissible when certain instructions are partially executed and then restarted:

- For any store, bytes at the target location may have been altered (if write access to that page in which bytes were altered is permitted by the access control mechanism). In addition, for store conditional instructions, **CR0** has been set to an undefined value, and it is undefined whether the reservation has been cleared or not.
- For any load, bytes at the addressed location may have been accessed (if read access to that page in which bytes were accessed is permitted by the access control mechanism).
- For load multiple or load string, some registers in the range to be loaded may have been altered. Including the addressing registers **rA** and possibly **rB** in the range to be loaded is a programming error, and thus the rules for partial execution do not protect these registers against overwriting.

In no case is access control violated.

As previously stated, elementary, aligned, guarded loads and stores are the only access instructions guaranteed not to be interrupted after being partially executed. The following list identifies the specific instruction types for which interruption after partial execution may occur, as well as the specific interrupt types that could cause the interruption:

1. Any load or store (except elementary, aligned, or guarded):
  - Any asynchronous interrupt
  - Machine check
  - Program (imprecise mode floating-point enabled)
  - Program (imprecise mode auxiliary processor enabled)
  - Decrementer
  - Fixed-interval timer
  - Watchdog timer
  - Debug (unconditional debug event)
2. Misaligned elementary load or store, or any multiple or string:
 

All of the above listed under item <sup>(1)</sup>, plus the following:

  - Alignment
  - Data storage (if the access crosses a protection boundary)
  - Debug (data address compare, data value compare)

The **mtcrf** and **mfcrr** instructions can also be partially executed due to the occurrence of any of the interrupts listed under item <sup>(1)</sup> at the time **mtcrf** or **mfcrr** was executing.

- All instructions before **mtcrf** or **mfcrr** have completed execution. Some memory accesses generated by these preceding instructions may not have completed.
- No subsequent instruction has begun execution.
- The **mtcrf** or **mfcrr** instruction, whose address was saved in **SRR0/CSRR0** at the time of the interrupt, may appear not to have begun or may have partially executed.

## 5.10 Interrupt ordering and masking

Multiple exceptions that can each generate an interrupt can exist simultaneously. However, the PowerPC architecture does not provide for reporting multiple simultaneous interrupts of the same class (critical or noncritical). Therefore, the PowerPC architecture defines that interrupts must be ordered with one another and provides a way to mask certain persistent interrupt types.

When an interrupt type is masked (disabled) and an event causes an exception that would normally generate an interrupt of that type, the exception persists as a status bit in a register (which register depends upon the exception type) but no interrupt is generated. Later, if the interrupt type is enabled (unmasked) and the exception status has not been cleared by software, the interrupt due to the original exception event is finally generated. (A typical implementation has such a mechanism for certain debug events. A signal that triggers an asynchronous interrupt, such as external input, must be asserted until they are taken. There is no mechanism for saving the external interrupt if the signal is negated before the interrupt is taken. All interrupts are level-sensitive except for machine check, which is edge-triggered.)

All asynchronous interrupt types and some synchronous interrupt types can be masked. An example of a maskable synchronous interrupt type is the floating-point enabled exception-type program interrupt. The execution of a floating-point instruction that causes FPSCR[FEX] to be set is considered an exception event, regardless of the setting of MSR[FE0,FE1]. If MSR[FE0,FE1] are both 0, the floating-point enabled exception-type program interrupt is masked, but the exception persists in FPSCR[FEX]. Later, if MSR[FE0,FE1] are enabled, the interrupt is generated.

The PowerPC architecture allows implementations to avoid situations in which an interrupt would cause state information (saved in save/restore registers) from a previous interrupt to be overwritten and lost. As a first step, upon any noncritical class interrupt, hardware automatically disables further asynchronous, noncritical class interrupts (external input) by clearing MSR[EE]. Likewise, upon any critical class interrupt, hardware automatically disables further asynchronous interrupts, both critical and noncritical, by clearing MSR[CE] and MSR[EE]. Critical input, watchdog timer, and debug interrupts are disabled by clearing MSR[CE,DE]. Note that machine check interrupts, while considered neither asynchronous nor synchronous, are not maskable by MSR[CE,DE,EE] and could be presented in a situation that could cause loss of state information.

This first step of clearing MSR[EE] (and MSR[CE,DE] for critical class interrupts) prevents subsequent asynchronous interrupts from overwriting save/restore registers before software can save their contents. On any interrupt, hardware also automatically clears MSR[WE,PR,FP,FE0,FE1,IS,DS], which helps avoid subsequent interrupts of certain other types. However, guaranteeing that these interrupt types do not occur also requires system software to avoid executing instructions that could cause (or enable) a subsequent interrupt, if SRR1 contents have not been saved.

### 5.10.1 Guidelines for system software

*Table 177* lists actions system software must avoid before saving save/restore register contents.

**Table 177. Operations to avoid**

Operation	Reason
Reenabling MSR[EE] (or MSR[CE,DE] in critical class interrupt handlers)	Prevents any asynchronous interrupts, and (in the case of MSR[DE]) any debug interrupts, including synchronous and asynchronous types
Branching (or sequential execution) to addresses not mapped by the TLB, mapped without UX = 1 or SX = 1 permission, or causing large address or instruction address overflow exceptions.	Prevents instruction storage, instruction TLB error, and instruction address overflow interrupts
Load, store, or cache management instructions to addresses not mapped by the TLB or not having required access permissions.	Prevents data storage and data TLB error interrupts
Execution of system call ( <b>sc</b> ) or trap ( <b>tw</b> , <b>twi</b> , <b>td</b> , <b>tdi</b> )	Prevents system call and trap exception-type program interrupts
Execution of any floating-point instruction	Prevents floating-point unavailable interrupts. Note that this interrupt would occur upon execution of any floating-point instruction, due to the automatic clearing of MSR[FP]. However, even if software were to reenabling MSR[FP], floating-point instructions must still be avoided to prevent program interrupts due to various possible program interrupt exceptions (floating-point enabled, unimplemented operation).
Reenabling of MSR[PR]	Prevents privileged instruction exception-type program interrupts. Alternatively, software could reenabling MSR[PR] but avoid executing any privileged instructions.
Execution of any auxiliary processor instruction	Prevents auxiliary processor unavailable, auxiliary processor enabled type, and unimplemented operation type program interrupts
Execution of any illegal instructions	Prevents illegal instruction exception-type program interrupts
Execution of any instruction that could cause an alignment interrupt	Prevents alignment interrupts, including string or multiple instructions and misaligned elementary load or store instructions. <a href="#">Section 5.7.6: Alignment interrupt</a> , lists instructions that cause alignment interrupts.

If the machine check APU is not implemented, machine check interrupts are a special case. Machine checks are critical interrupts, but normal critical interrupts (critical input, watchdog timer, and debug) do not automatically disable machine checks. Machine checks are disabled by clearing MSR[ME], and only a machine check interrupt itself automatically clears this bit. Thus there is always the risk that a machine check interrupt could occur within a normal critical interrupt handler before it saves the save/restore registers' contents. In such a case, the interrupt may not be recoverable.



It is unnecessary for hardware or software to avoid critical-class interrupts from within noncritical-class interrupt handlers (hence hardware does not automatically clear MSR[CE,ME,DE] on a noncritical interrupt), since the two interrupt classes use different save/restore registers. However, because a critical class interrupt can occur within a noncritical handler before the noncritical handler saves SRR0/SRR1, hardware and software must cooperate to avoid both critical and noncritical class interrupts from within critical class interrupt handlers. Therefore, within the critical class interrupt handler, both pairs of save/restore registers may contain data necessary to system software.

### 5.10.2 Interrupt order

Enabled interrupt types for which simultaneous exceptions can exist are prioritized as follows:

1. Synchronous (non-debug) interrupts:
  - Data storage
  - Instruction storage
  - Alignment
  - Program
  - Floating-point unit unavailable
  - Auxiliary processor unavailable
  - System call
  - Data TLB error
  - Instruction TLB error

Only one of the above synchronous interrupt types may have an existing exception generating it at a given time. This is guaranteed by the exception priority mechanism (see [Section 5.11: Exception priorities](#)) and the sequential execution model.

2. Machine check
3. Debug
4. Critical input
5. Watchdog timer
6. External input
7. Fixed-interval timer
8. Decrementer

Although, as indicated above, noncritical, synchronous exception types listed under item <sup>(1)</sup> are generated with higher priority than critical interrupt types in items <sup>(2)</sup>–<sup>(5)</sup>, noncritical interrupts are immediately followed by the highest priority existing critical interrupt type, without executing any instructions at the noncritical interrupt handler. This is because noncritical interrupt types do not automatically disable MSR mask bits for critical interrupt types (CE and ME). In all other cases, a particular interrupt type listed above automatically disables subsequent interrupts of the same type, as well as all lower priority interrupt types.

## 5.11 Exception priorities

Book E requires all synchronous (precise and imprecise) exceptions to be reported in program order, as required by the sequential execution model. The one exception to this rule is the case of multiple synchronous imprecise exceptions. Upon a synchronizing event,

all previously executed instructions are required to report any synchronous imprecise interrupt-generating exceptions, and the interrupt is then generated with all of those exception types reported cumulatively in the ESR and in any status registers associated with the particular exception type (such as the FPSCR).

For any single instruction attempting to cause multiple exceptions for which the corresponding synchronous interrupt types are enabled, this section defines the priority order by which the instruction is permitted to cause a single enabled exception, thus generating a particular synchronous interrupt. Note that it is this exception priority mechanism, along with the requirement that synchronous interrupts be generated in program order, that guarantees that at any given time there exists for consideration only one of the synchronous interrupt types listed in item 1 of [Section 5.10.2: Interrupt order](#). The exception priority mechanism also prevents certain debug exceptions from existing in combination with certain other synchronous interrupt-generating exceptions.

The EIS defines priorities for all exceptions including those defined in optional APUs. Interrupt types are defined as either synchronous (the interrupt is as a direct result of an instruction in execution) or asynchronous, (the interrupt results from an event external to the execution of a particular instruction or an instruction removes a gating condition to a pending exception). Except for machine check interrupts, which can be either synchronous or asynchronous, interrupts are either synchronous or asynchronous exclusively.

Because asynchronous interrupts may temporally be sampled either before or after an instruction is completed, an implementation can order asynchronous interrupts among only asynchronous interrupts and order synchronous interrupts among only synchronous interrupt. The distinction is important because synchronous interrupts that require post-completion actions (such as system call or debug instruction complete exceptions) cannot be separated from the completion of the instruction. Therefore, asynchronous interrupts cannot be sampled during the completion and post-completion synchronous exceptions for a given instruction.

The relative priorities for asynchronous exceptions is given in [Table 178](#) and for synchronous exceptions in and [Table 179](#). In many cases, certain exceptions cannot occur at the same time (for example, program-trap and program-Illegal cannot occur on the same instruction). In general those exceptions are grouped at the same relative priority.

**Table 178. EIS asynchronous exception priorities**

Relative priority	Exception	Interrupt level <sup>(1)</sup>	Interrupt nature	Pre/post completion <sup>(2)</sup>	Comments
0	Machine check	Machine check	Asynch/synch	pre for synch	Asynchronous exceptions may come from processor or from an external source.

**Table 178. EIS asynchronous exception priorities (continued)**

Relative priority	Exception	Interrupt level <sup>(1)</sup>	Interrupt nature	Pre/post completion <sup>(2)</sup>	Comments
1	Debug—UDE	Critical/debug	Asynch	N/A	Generally used for an externally generated high priority attention signal.
	Debug—IDE	Critical/debug	Asynch	N/A	Usually taken after MSR[DE] goes from 0 to 1 via <b>rfdi/rfci</b> or <b>mtmsr</b> .
	Debug—interrupt taken	Critical/debug	Asynch	N/A	Debug interrupt taken after original interrupt changed NIA and MSR.
	Debug—critical interrupt taken	Debug	Asynch	N/A	Debug interrupt taken after original critical interrupt has changed NIA and MSR.
2	Critical input	Critical	Asynch	N/A	
3	Watchdog	Critical	Asynch	N/A	
4	External input	Base	Asynch	N/A	
18	Fixed interval timer	Base	Asynch	N/A	
19	Decrementer	Base	Asynch	N/A	
20	Performance Monitor	Base	Asynch	N/A	

1. The interrupt level defines the set of save/restore registers used when the interrupt is taken—base (SRR0/SRR1), critical (CSRR0/CSRR1), debug (DSRR0/DSRR1), and machine check (MCSR0/MCSR1).
2. Pre- or post-completion refers to whether the exception occurs before an instruction completes (pre) and the corresponding interrupt points to the instruction causing the exception, or if the instruction completes (post) and the corresponding interrupt points to the next instruction to be executed.

**Table 179. EIS synchronous exception priorities**

Relative priority	Exception	Interrupt level <sup>(1)</sup>	Interrupt nature	Pre/post completion <sup>(2)</sup>	Comments
5	Debug—instruction address compare	Critical/debug	Synch	pre	
6	ITLB	Base	Synch	pre	
	ISI	Base	Synch	pre	

**Table 179. EIS synchronous exception priorities (continued)**

Relative priority	Exception	Interrupt level <sup>(1)</sup>	Interrupt nature	Pre/post completion <sup>(2)</sup>	Comments
7	FP unavailable	Base	Synch	pre	
	AltiVec unavailable	Base	Synch	pre	Defined by the AltiVec APU.
	SPE unavailable	Base	Synch	pre	Defined by the SPE APU.
	Embedded floating-point unavailable	Base	Synch	pre	Defined by the embedded floating point APUs.
8	Debug—trap	Critical/debug	Synch	pre	
9	Program—illegal instruction	Base	Synch	pre	
	Program—unimplemented operation	Base	Synch	pre	
	Program—privileged instruction	Base	Synch	pre	
	Program—Trap	Base	Synch	pre	
	Program—FP enabled	Base	Synch	pre	An FP enabled interrupt may be imprecise.
10	(Alignment)	Base	Synch	pre	Alignment may be handled at either priority.
11	DTLB	Base	Synch	pre	
	Data storage	Base	Synch	pre	
12	Alignment	Base	Synch	pre	Alignment may be handled at either priority.
13	System call	Base	Synch	post	Points SRR0 to instruction after <b>sc</b> (post completion).
	Embedded FP data	Base	Synch	pre	Defined by the SPE APU.
	Embedded FP round	Base	Synch	post	Points SRR0 to the instruction after the one causing the exception (post completion). Defined by SPE APU.
	AltiVec Assist	Base	Synch	pre	Defined by the AltiVec APU.

**Table 179. EIS synchronous exception priorities (continued)**

Relative priority	Exception	Interrupt level <sup>(1)</sup>	Interrupt nature	Pre/post completion <sup>(2)</sup>	Comments
14	Debug—return from interrupt	Critical/debug	Synch	pre	
	Debug—Return from critical interrupt	Debug	Synch	pre	Defined by the enhanced debug APU.
	Debug—branch taken	Critical/debug	Synch	pre	
15	Debug—DAC	Critical/debug	Synch	pre or post	Preferred method is pre-completion.
16	Debug—DVC	Critical/debug	Synch	pre or post	Preferred method is pre-completion.
17	Debug—instruction complete	Critical/debug	Synch	post	Points [CD]SRR0 to next instruction (post completion).

1. The interrupt level defines the set of save/restore registers used when the interrupt is taken—base (SRR0/SRR1), critical (CSRR0/CSRR1), debug (DSRR0/DSRR1), and machine check (MCSRR0/MCSRR1).
2. Pre- or post-completion refers to whether the exception occurs before an instruction completes (pre) and the corresponding interrupt points to the instruction causing the exception, or if the instruction completes (post) and the corresponding interrupt points to the next instruction to be executed.

## 6 Storage architecture

This chapter describes the cache and MMU portions of the Book E implementation standards (EIS). Note that not all features that are defined by the EIS storage architecture are supported on all ST EIS processors; consult the user documentation. This chapter is organized into three sections:

- [Section 6.2: Memory and cache coherency](#)
- [Section 6.3: Cache model](#)
- [Section 6.4: Storage model](#)

### 6.1 Overview

The Book E architecture memory and cache definitions support a wide variety of embedded implementations. To provide such flexibility, Book E defines many features in a very general way, leaving specific details up to the implementation. To ensure consistency among its Book E cores and devices, ST has defined more specific implementation standards. However, these standards still leave many details up to individual implementations. To provide context for those features, this chapter describes aspects of the memory hierarchy and the memory management model defined by Book E; it also describes the ST EIS.

*Note: This chapter describes some features (in particular, registers) in a very general way that does not include some details that are important to the programmer. There are also small differences in how some features are defined here and how they are implemented. For implementation-specific details, see the user documentation.*

*Throughout this chapter, references to load instructions include cache management and other instructions that are stated in the instruction descriptions to be treated as a load, and references to store instructions include the cache management and other instructions that are treated as a store.*

The following APUs, which are part of the EIS storage architecture, are defined in [Section 9: Storage-related APUs](#):

- Cache line locking APU
- Cache way partitioning APU
- Direct cache flush APU

These APUs may be implemented independently of each other. They are defined together in a single specification because it is likely that an implementation will include more than one of these APUs.

### 6.2 Memory and cache coherency

The primary objective of a coherent memory system is to provide the same image of memory to all devices using the system. Coherency allows synchronization and cooperative use of shared resources. Otherwise, multiple copies of data corresponding to a memory location, some containing outdated values, could exist in a system, resulting in errors when the outdated values are used. Each memory-sharing device must follow rules for managing the state of its cache. This section describes the coherency mechanisms of the Book E architecture and the cache coherency protocols that the ST Book E devices support.

Unless specifically noted, the discussion of coherency in this section applies to the core complex data cache only. The instruction cache is not snooped for general coherency with other caches; however, it is snooped when the Instruction Cache Block Invalidate (**icbi**) instruction is executed by this processor or any processor in the system.

## 6.2.1 Memory/Cache access attributes

Some memory characteristics can be set on a page basis by using the WIMGE bits in the translation lookaside buffer (TLB) entries. These bits allow both uniprocessor and multiprocessor system designs to exploit numerous system-level performance optimizations. The WIMGE attributes control the following:

- Write-through (W bit)
- Caching-inhibited (I bit)
- Memory-coherency-required (M bit)
- Guarded (G bit)
- Endianness (E bit)

In addition to the WIMGE bits, the Book E MMU model defines the following attributes on a page basis:

- User-definable (U0, U1, U2, U3)

The EIS defines the following optional attributes, which are manipulated by software through MMU assist register 2 (MAS2):

- Alternate coherency mode (ACM). The ACM attribute, programmed through MAS2[ACM], allows an implementation to employ multiple coherency methods and to participate in multiple coherency protocols. If the M attribute (memory coherence required) is not set for a page (M = 0), the page has no coherency associated with it and the ACM attribute is ignored. If the M attribute is set for a page (M = 1), the ACM attribute determines the coherency domain (or protocol) used. ACM values are implementation dependent.
- Variable length encoding (VLE). The VLE attribute, MAS2[VLE], identifies pages that contain instructions from the VLE instruction set. If VLE = 0, instructions fetched from the page are decoded and executed as PowerPC (and associated EIS APUs) instructions. If VLE = 1, instructions fetched from the page are decoded and executed as Power Embedded instructions.

Consult the user documentation to determine whether the EIS-defined attributes are implemented.

The WIMGE attributes are programmed by the operating system for each page. The W and I attributes control how the processor performing an access uses its own cache. The M attribute ensures that coherency is maintained for all copies of the addressed memory location. The G attribute prevents speculative loading from the addressed memory location. (An operation is said to be performed speculatively if, at the time that it is performed, it is not known to be required by the sequential execution model.) The E attribute defines the order in which the bytes that comprise a multiple-byte data object are stored in memory (big- or little-endian).

The WIMGE attributes occupy 5 bits in the TLB entries for page address translation. The operating system writes the WIMGE bits for each page into the TLB entries in system memory as it maps translations. For more information, see [Section 6.4.10.2: TLB entries](#).

All combinations of these attributes are supported except those that simultaneously specify a region as write-through and caching-inhibited. Write-through and caching-inhibited attributes are mutually exclusive because the write-through attribute permits the data to be in the data cache while the caching-inhibited attribute does not.

Memory that is write-through or caching-inhibited is not intended for general-purpose programming. For example, **lwarx** and **stwcx**. instructions may cause the system DSI exception handler to be invoked if they specify a location in memory having either of these attributes. Some implementations take a data storage interrupt if the location is write-through but does not take the interrupt if the location is cache-inhibited. Note that, except that the guarded bit does not prevent instruction prefetches, the definitions of the WIMG bits are unchanged

### 6.2.1.1 Write-through attribute

A page marked  $W = 0$  is considered to be write-back. If some store instructions executed by a given processor access locations in a block as write-through and other store instructions executed by the same processor access locations in that block as write-back, software must ensure that the block cannot be accessed by another processor or mechanism in the system.

A store to a write-through ( $W = 1$ ) memory location is performed in main memory and may cause additional memory locations to be accessed. If a copy of the block containing the specified location is retained in the data cache, the store is also performed in the data cache. A store to write-through memory cannot cause a block to be put in a modified state in the data cache.

Also, if a store instruction that accesses a block in a location marked as write-through is executed when the block is already considered to be modified in the data cache, the block may continue to be considered to be modified in the data cache even if the store causes all modified locations in the block to be written to main memory. In some processors, accesses caused by separate store instructions that specify locations in write-through memory may be combined into one access. This is called store-gathering. Such combining does not occur if the store instructions are separated by an **msync** or an **mbar**.

### 6.2.1.2 Caching-inhibited attribute

A load instruction that specifies a location in caching-inhibited ( $I = 1$ ) memory is performed to main memory and may cause additional locations in main memory to be accessed unless the specified location is also guarded. An instruction fetch from caching-inhibited memory may cause additional words in main memory to be accessed. No copy of the accessed locations is placed into the caches.

In some processors, nonoverlapping accesses caused by separate load instructions that specify locations in caching-inhibited memory may be combined into one access, as may nonoverlapping accesses caused by separate store instructions to caching-inhibited memory (that is, store-gathering). Such combining does not occur if the load or store instructions are separated by an **msync** instruction, or by an **mbar** instruction if the memory is also guarded.

### 6.2.1.3 Memory-coherence-required attribute

Memory coherence refers to the ordering of stores to a single location. Atomic stores to a given location are coherent if they are serialized in some order, and no processor or mechanism is able to observe any subset of those stores as occurring in a conflicting order.



This serialization order is an abstract sequence of values; the physical location need not assume each of the values written to it. For example, a processor may update a location several times before the value is written to physical memory.

The result of a store operation is not available to every processor or mechanism at the same instant, and it may be that a processor or mechanism observes only some of the values that are written to a location. However, when a location is accessed atomically and coherently by all processors and mechanisms, the sequence of values loaded from the location by any processor or mechanism during any interval of time forms a sub-sequence of the sequence of values that the location logically held during that interval. That is, a processor or mechanism can never load a newer value first and then, later, load an older value.

Memory coherence is managed in blocks called coherence blocks. Although a block's size is implementation-dependent, it is usually larger than a word and is often the size of a cache block.

When memory coherence is not required ( $M = 0$ ), the hardware need not enforce data coherence for memory accesses initiated by the processor. When memory coherence is required ( $M = 1$ ), the hardware must enforce data coherence for memory accesses initiated by the processor. Hardware support for the memory-coherence-required attribute is optional for implementations that do not support multiprocessing.

#### 6.2.1.4 Guarded attribute

When the guarded bit is set, the page is designated as guarded. This setting can be used to protect certain memory areas from read accesses made by the processor that are not dictated directly by the program. If areas of physical memory are not fully populated (in other words, there are holes in the physical memory map within this area), this setting can protect the system from undesired accesses caused by speculative (referred to as 'out of order' in the architecture specification, and described in [Section 6.2.1.5: Definition of a speculative and out-of-order memory accesses](#)) load operations that could lead to the generation of the machine check exception. Also, the guarded bit can be used to prevent speculative load operations from occurring to certain peripheral devices that produce undesired results when accessed in this way.

#### 6.2.1.5 Definition of a speculative and out-of-order memory accesses

In the architecture definition, the term 'out of order' replaced the term 'speculative' with respect to memory accesses to avoid a conflict between the word's meaning in the context of execution of instructions past unresolved branches. The architecture's use of out of order in this context could in turn be confused with the notion of loads and stores being reordered in a weakly ordered memory system.

In the context of memory accesses, this document uses the terms 'speculative' and 'out of order' as follows:

- Speculative memory access—An access to memory that occurs before it is known to be required by the sequential execution model.
- Out-of-order memory access—A memory access performed ahead of one that may have preceded it in the sequential model, such as is allowed by a weakly ordered memory model.

### 6.2.1.6 Performing operations speculatively

An operation is said to be nonspeculative if it is guaranteed to be required by the sequential execution model. Any other operation is said to be performed speculatively, which the architecture specification refers to as out of order.

Operations are performed speculatively by hardware on the expectation that the results will be needed by an instruction that will be required by the sequential execution model. Whether the results are needed depends on whether control flow is diverted away from the instruction by an event such as an exception, branch, trap, system call, return from interrupt instruction, or anything else that changes the context in which the instruction is executed.

Typically, the hardware performs operations speculatively when it has resources that would otherwise be idle, so the operation incurs little or no cost. If subsequent events such as branches or exceptions indicate that the operation would not have been performed, the processor abandons any results of the operation except as described below.

Most operations can be performed speculatively, as long as the machine appears to follow the sequential execution model. Certain speculative operations are restricted, as follows:

- Stores—A store instruction cannot execute speculatively in a manner such that the alteration of the target location can be observed by other processors or mechanisms.
- Accessing guarded memory—The restrictions for this case are given in [Section 6.2.1.7.1: Speculative accesses to guarded memory](#).

No error of any kind other than a machine check exception may be reported due to an operation that is performed speculatively, until such time as it is known that the operation is required by the sequential execution model. The only other permitted side effect (other than machine check) of performing an operation speculatively is that nonguarded memory locations that could be fetched into a cache by nonspeculative execution may be fetched speculatively into that cache.

### 6.2.1.7 Guarded memory

Memory is said to be well behaved if the corresponding physical memory exists and is not defective, and if the effects of a single access to it are indistinguishable from the effects of multiple identical accesses to it. Data and instructions can be fetched speculatively from well-behaved memory without causing undesired side effects.

Memory is said to be guarded if the G bit is set for the page. In general, memory that is not well-behaved should be guarded. Because such memory may represent an I/O device or include nonexistent locations, a speculative access to such memory may cause an I/O device to perform incorrect operations or may cause a machine check.

Note that if separate store instructions access memory that is both caching-inhibited and guarded, the accesses are performed in the order specified by the program. If an aligned load or store that is not a string or multiple access to caching-inhibited, guarded memory has accessed main memory and an external, decremter, or imprecise-mode floating-point enabled exception is pending, the load or store is completed before the exception is taken.

#### 6.2.1.7.1 Speculative accesses to guarded memory

Accesses for load instructions from guarded memory may be performed speculatively if a copy of the target location is in a cache; in this case, the location may be accessed from the cache or from main memory.

Note that software should ensure that only well-behaved memory is loaded into a cache, either by marking as caching-inhibited (and guarded) all memory that may not be well-behaved or by marking such memory caching-allowed (and guarded) and referring only to cache blocks that are well-behaved.

Instruction accesses: guarded memory and no-execute memory

The G bit is ignored for instruction fetches, and instructions are speculatively fetched from guarded pages. To prevent speculative fetches from pages that do not contain instructions and are not well-behaved, the page should be designated as no-execute (with the UX/SX page permission bits cleared). If the effective address of the current instruction is mapped to no-execute memory, an ISI exception is generated.

#### 6.2.1.7.2 Endianness

Objects may be loaded from or stored to memory in byte, half-word, word, or double-word units. For a particular data length, the load and store operations are symmetrical; a store followed by a load of the same data object yields an unchanged value. Book E makes no guarantees about the order in which the bytes that comprise multiple-byte data objects are stored into memory. The endianness (E) page attribute distinguishes between memory that is big or little endian, as described in the following subsections.

Except for instruction fetches, it is always permitted to access the same location using two effective addresses with different E bit settings. Instruction pages must be flushed from any caches before the E bit can be changed for those addresses. See [Section 4.2.3.4: Byte ordering](#), for more information about endianness.

#### 6.2.1.7.3 Big-endian pages

If a stored multiple-byte object is probed by reading its component bytes one at a time using load-byte instructions, the store order may be perceived. If such probing shows that the lowest memory address contains the highest-order byte of the multiple-byte scalar, the next-higher sequential address the next-least-significant byte, and so on, the multiple-byte object is stored in big-endian form. Big-endian memory is defined on a page basis by the memory/cache attribute, E = 0.

Note that strings are not multiple-byte scalars but are interpreted as a series of single-byte scalars. Bytes in a string are loaded from memory using a load string word instruction, starting at the lowest-numbered address, and placed into the target register or registers starting at the left-most byte of the least-significant word. Bytes in a string are stored using a store string word instruction from the source register, starting at the left-most byte of the least-significant word, and placed into memory, starting at the lowest-numbered address.

#### 6.2.1.7.4 Little-endian pages

Alternatively, if the probing shows that the lowest memory address contains the lowest-order byte of the multiple-byte scalar, the next-higher sequential address the next-most-significant byte, and so on, the multiple-byte object is stored in little-endian form. Little-endian memory is defined on a page basis by the memory/cache attribute, E = 1, and for Book E devices is defined as true little-endian memory.

#### 6.2.1.7.5 Structure mapping examples

The following C programming example defines the data structure S used in this section to demonstrate how the bytes that comprise each element (a, b, c, d, e, and f) are mapped into

memory. The structure contains scalars (shown in hexadecimal in the comments) and a sequence of characters, shown in single quotation marks.

```

struct {
    int      a;          /* 0x1112_1314      word*/
    double   b;          /* 0x2122_2324_2526_2728double word*/
    char *   c;          /* 0x3132_3334      word*/
    char     d[7];       /* 'L','M','N','O','P','Q','R' array of bytes*/
    short    e;          /* 0x5152           half word*/
    int      f;          /* 0x6162_6364      word*/
} S;
    
```

big-endian mapping of the structure ia shown below.

**6.2.1.7.6 Big-endian mapping of structure**

Contents	11	12	13	14	(x)	(x)	(x)	(x)
Address	00	01	02	03	04	05	06	07
Contents	21	22	23	24	25	26	27	28
Address	08	09	0A	0B	0C	0D	0E	0F
Contents	31	32	33	34	'L'	'M'	'N'	'O'
Address	10	11	12	13	14	15	16	17
Contents	'P'	'Q'	'R'	(x)	51	52	(x)	(x)
Address	18	19	1A	1B	1C	1D	1E	1F
Contents	61	62	63	64	(x)	(x)	(x)	(x)
Address	20	21	22	23	24	25	26	27

Note that the MSB of each scalar is at the lowest address. The mapping uses padding (indicated by (x)) to align the scalars—4 bytes between elements a and b, 1 byte between d and e, and 2 bytes between e and f. Note that the padding is determined by the compiler, not the architecture.

The structure using little-endian mapping, showing double words laid out with addresses increasing from right to left.

Little-endian mapping of structure S—alternate view

Contents	(x)	(x)	(x)	(x)	11	12	13	14
Address	07	06	05	04	03	02	01	00
Contents	21	22	23	24	25	26	27	28
Address	0F	0E	0D	0C	0B	0A	09	08
Contents	'O'	'N'	'M'	'L'	31	32	33	34
Address	17	16	15	14	13	12	11	10
Contents	(x)	(x)	51	52	(x)	'R'	'Q'	'P'
Address	1F	1E	1D	1C	1B	1A	19	18
Contents	(x)	(x)	(x)	(x)	61	62	63	64
Address	27	26	25	24	23	22	21	20

### 6.2.1.8 Mismatched memory cache attributes

Accesses to the same memory location using two effective addresses for which the write-through required attribute (W bit) differs meet the memory coherence requirements described in [Section 6.2.1.1: Write-through attribute](#), if the accesses are performed by a single processor. If the accesses are performed by two or more processors, coherence is enforced by the hardware only if the write-through attribute is the same for all the accesses.

Loads, stores, **dcbz** instructions, and instruction fetches to the same memory location using two effective addresses for which the caching-inhibited attribute (I bit) differs must meet the requirement that a copy of the target location of an access to caching-inhibited memory not be in the cache. Violation of this requirement is considered a programming error; software must ensure that the location has not previously been brought into the cache or, if it has, that it has been flushed from the cache. If the programming error occurs, the result of the access is boundedly undefined. It is not considered a programming error if the target location of any other cache management instruction to caching-inhibited memory is in the cache.

Accesses to the same memory location using two effective addresses for which the memory coherence attribute (M bit) differs may require explicit software synchronization before accessing the location with  $M = 1$  if the location has previously been accessed with  $M = 0$ . Any such requirement is system-dependent. For example, in some systems that use bus snooping, no software synchronization may be required. In some directory-based systems, software may be required to execute **dcbf** instructions on each processor to flush all cache entries accessed with  $M = 0$  before accessing those locations with  $M = 1$ .

Accesses to the same memory location using two effective addresses for which the guarded attribute (G bit) differs are always permitted.

Except for instruction fetches, accesses to the same memory location using two effective addresses for which the endian storage attribute (E bit) differs are always permitted as described in [Section 6.2.1.7.2: Endianness](#). Instruction memory locations must be flushed before the endian attribute can be changed for those addresses.

The requirements on mismatched user-defined memory attributes (U0–U3) is implementation-dependent.

### 6.2.1.9 Coherency paradoxes and WIMGE

Care must be taken with respect to the use of the WIMGE bits if coherent memory support is desired. Careless programming of these bits may create situations that present coherency paradoxes to the processor. These paradoxes can occur within a single processor or across several processors. It is important to note that, in the presence of a paradox, the operating system software is responsible for correctness.

In particular, a coherency paradox can occur when the state of these bits is changed without appropriate precautions (such as flushing the pages that correspond to the changed bits from the caches of all processors in the system) or when the address translations of aliased real addresses specify different values for certain WIMGE bit values. For more information, see [Section 6.2.1.8: Mismatched memory cache attributes](#).

Support for M = 1 memory is optional. Cache attribute settings where both W = 1 and I = 1 are not supported. For all supported combinations of the W, I, and M bits, both G and E may be 0 or 1.

The default setting of the WIMGE bits is 0b01000.

### 6.2.1.10 Self-modifying code

When a processor modifies any memory location that can contain an instruction, **software must ensure that** the instruction cache is made consistent with data memory and that the modifications are made visible to the instruction fetching mechanism. This must be done even if the cache is disabled or if the page is marked caching-inhibited.

The following instruction sequence can be used to accomplish this when the instructions being modified are in memory that is memory-coherence required and one processor both modifies the instructions and executes them. (Additional synchronization is needed when one processor modifies instructions that another will execute.)

The following sequence synchronizes the instruction stream (using either **dcbst** or **dcbf**):

<b>dcbst</b> (or <b>dcbf</b> )	update memory
<b>msync</b>	wait for update
<b>icbi</b>	remove (invalidate) copy from instruction cache
<b>msync</b>	ensure that ICBI invalidation at icache has completed
<b>isync</b>	remove copy in own instruction buffer

## 6.2.2 Shared memory

The architecture supports sharing memory between programs, between different instances of the same program, and between processors and other mechanisms. It also supports access to a memory location by one or more programs using different effective addresses. In these cases, memory is shared in blocks that are an integral number of pages. When one physical memory location has different effective addresses, the addresses are said to be aliases. Each application can be granted separate access privileges to aliased pages.

[Section 6.2.2.2: Lock acquisition and import barriers](#), gives examples of how **msync** and **mbar** are used to control memory access ordering when memory is shared among programs.

### 6.2.2.1 Memory access ordering

The memory model in Book E for memory access ordering is weakly consistent. This provides an opportunity for improved performance over a model with stronger consistency

rules but places the responsibility on the program to ensure that ordering or synchronization instructions are properly placed for correct execution of the program.

The order in which a processor accesses memory, the order in which those accesses are performed with respect to other processors or mechanisms, and the order in which they are performed in main memory may all be different. [Table 180](#) describes how the architecture defines requirements for ordering of loads and stores.

**Table 180. Load and store ordering**

Type of Access	Architecture definition
Load ordering with respect to other loads	The architecture guarantees that loads that are both caching-inhibited (I = 1) and guarded (G = 1) are not reordered with respect to one another.
	If a load instruction depends on the value returned by a preceding load (because the value is used to compute the effective address specified by the second load), the corresponding memory accesses are performed in program order with respect to any processor or mechanism to the extent required by the associated memory coherence required attributes (that is, the memory coherence required attribute, if any, associated with each access). This applies even if the dependency does not affect program logic (for example, the value returned by the first load is ANDed with zero and then added to the effective address specified by the second load).
Store ordering with respect to other stores	If two store instructions specify memory locations that are both caching inhibited and guarded, the corresponding memory accesses are performed in program order with respect to any processor or mechanism. Otherwise, stores are weakly ordered with respect to one another.
Store ordering with respect to loads	The architecture specifies that an <b>msync</b> or <b>mbar</b> must be used to ensure sequential ordering of loads with respect to stores.

When a processor (P1) executes **msync** or **mbar**, a memory barrier is created that separates applicable memory accesses into two groups, G1 and G2. G1 includes all applicable memory accesses associated with instructions preceding the barrier-creating instruction, and G2 includes all applicable memory accesses associated with instructions following the barrier-creating instruction.

[Table 181](#) shows an example using a two-processor system.

**Table 181. Memory barrier when coherency is required (M = 1)**

Processor 1 (P1)	Memory access groups G1 and G2	Processor 2 (P2)
Instruction 1	G1: Memory accesses generated by P1 before the memory barrier	When memory coherence is required, G1 accesses that affect P2 are also performed before the memory barrier.
Instruction 2		
Instruction 3		
Instruction 4		
Instruction 5 ( <b>msync</b> or <b>mbar</b> )—Memory barrier		Barrier generated by P1 does not order P2 instructions or associated accesses with respect to other P2 instructions and associated accesses.

**Table 181. Memory barrier when coherency is required (M = 1) (continued)**

Processor 1 (P1)	Memory access groups G1 and G2	Processor 2 (P2)
Instruction 6	G2: Memory accesses generated by P1 after the memory barrier	When memory coherence is required, G2 accesses that affect P2 are also performed after the memory barrier.
Instruction 7		
Instruction 8		
Instruction 9		
Instruction 10		

The memory barrier ensures that all memory accesses in G1 are performed with respect to any processor or mechanism, to the extent required by the associated memory coherence required attributes (that is, the memory-coherence required attribute, if any, associated with each access), before any memory accesses in G2 are performed with respect to that processor or mechanism.

The ordering enforced by a memory barrier is said to be cumulative if it also orders memory accesses that are performed by processors and mechanisms other than P1, as follows:

- G1 includes all applicable memory accesses by any such processor or mechanism that have been performed with respect to P1 before the memory barrier is created.
- G2 includes all applicable memory accesses by any such processor or mechanism that are performed after a load instruction executed by that processor or mechanism has returned the value stored by a store that is in G2.

Table 182 shows an example of a cumulative memory barrier in a two-processor system.

**Table 182. Cumulative memory barrier**

Processor 1 (P1)	Memory access groups G1 and G2	Processor 2 (P2)
P1 Instruction 1	G1: Memory accesses generated by P1 and P2 that affect P1. Includes accesses generated by executing P2 instructions L–O (assuming that the access generated by instruction O occurs before P1’s <b>msync</b> is executed).	P2 Instruction L
P1 Instruction 2		P2 Instruction M
P1 Instruction 3		P2 Instruction N
P1 Instruction 4		P2 Instruction O
P1 Instruction 5 ( <b>msync</b> )—Cumulative memory barrier applies to all accesses except those associated with fetching instructions following <b>msync</b> .		P2 Instruction P
		P2 Instruction Q
		P2 Instruction R
P1 Instruction 6	G2: Memory accesses generated by P1 and P2. Includes accesses generated by P2 instructions P–X (assuming that the access generated by instruction P occurs after P1’s <b>msync</b> is executed) performed after a load instruction executed by P2 has returned the value stored by a store that is in G2.	P2 Instruction S
P1 Instruction 7		P2 Instruction T
P1 Instruction 8		P2 Instruction U
P1 Instruction 9		P2 Instruction V
P1 Instruction 10		The <b>msync</b> memory barrier does not affect accesses associated with instruction fetching that occur after the <b>msync</b> .
P1 Instruction 11	P2 Instruction X	



A memory barrier created by **msync** is cumulative and applies to all accesses except those associated with fetching instructions following the **msync**. See the definition of **mbar** in [Section 4.3.1.15: Memory synchronization instructions](#), for a description of the corresponding properties of the memory barrier created by that instruction.

### 6.2.2.1.1 Programming considerations

Because stores cannot be performed out of program order, as described in Book E, if a store instruction depends on the value returned by a preceding load (because the value the load returns is needed to compute either the effective address specified by the store or the value to be stored), the corresponding accesses are guaranteed to be performed in program order. The same applies whether or not the store instruction executes is dependent upon a conditional branch that in turn depends on the value returned by a preceding load. For example, if a conditional branch depends on a preceding load and that branch chooses between a path that includes a store instruction if the condition is met, that dependent store is not performed unless and until the condition determined by the load is met.

Because instructions following an **isync** cannot execute until all instructions preceding **isync** have completed, if an **isync** follows a conditional branch instruction that depends on the value returned by a preceding load instruction, that load is performed before any loads caused by instructions following the **isync**. This is true even if the effects of the dependency are independent of the value loaded (for example, the value is compared to itself and the branch tests  $CRn[EQ]$ ), and even if the branch target is the next sequential instruction.

Except for the cases described above and earlier in this section, data and control dependencies do not order memory accesses. Examples include the following:

- If a load specifies the same memory location as a preceding store and the location is not caching inhibited, the load may be satisfied from a store queue (a buffer into which the processor places stored values before presenting them to the memory subsystem) and not be visible to other processors and mechanisms. As a result, if a subsequent store depends on the value returned by the load, the two stores need not be performed in program order with respect to other processors and mechanisms.
- Because a store conditional instruction may complete before its store is performed, a conditional branch instruction that depends on the CR0 value set by a store conditional instruction does not order that store with respect to memory accesses caused by instructions that follow the branch.

For example, in the following sequence, the **stw** is the **bc** instruction's target:

```

stwcx.
bc
stw

```

To complete, the **stwcx.** must update the architected CR0 value, even though its store may not have been performed. The architecture does not require that the store generated by the **stwcx.** must be performed before the store generated by the **stw**.

- Because processors may predict branch target addresses and branch condition resolution, control dependencies (branches, for example) do not order memory accesses except as described above. For example, when a subroutine returns to its caller, the return address may be predicted, with the result that loads caused by instructions at or after the return address may be performed before the load that obtains the return address is performed.

Some processors implement nonarchitected duplicates of architected resources such as GPRs, CR fields, and the LR, so resource dependencies (for example, specification of the same target register for two load instructions) do not force ordering of memory accesses.

Examples of correct uses of dependencies, **msync**, and **mbar** to order memory accesses can be found in hi.

Because the memory model is weakly consistent, the sequential execution model as applied to instructions that cause memory accesses guarantees only that those accesses appear to be performed in program order with respect to the processor executing the instructions. For example, an instruction may complete, and subsequent instructions may be executed, before memory accesses caused by the first instruction have been performed. However, for a sequence of atomic accesses to the same memory location for which memory coherence is required, the definition of coherence guarantees that the accesses are performed in program order with respect to any processor or mechanism that accesses the location coherently, and similarly if the location is one for which caching is inhibited.

Because caching-inhibited memory accesses are performed in main memory, memory barriers and dependencies on load instructions order such accesses with respect to any processor or mechanism even if the memory is not marked as requiring memory coherence.

#### 6.2.2.1.2 Programming examples

Example 1 shows cumulative ordering of memory accesses preceding a memory barrier, Example 2 shows cumulative ordering of memory accesses following a memory barrier. In both examples, assume that locations X, Y, and Z initially contain the value 0. In both, cumulative ordering dictates that the value loaded from location X by processor C is 1.

##### Example 1:

- Processor A stores the value 1 to location X.
- Processor B loads from location X obtaining the value 1, executes an **msync**, then stores the value 2 to location Y.
- Processor C loads from location Y obtaining the value 2, executes an **msync**, then loads from location X.

##### Example 2:

- Processor A stores the value 1 to location X, executes an **msync**, then stores the value 2 to location Y.
- Processor B loops, loading from location Y until the value 2 is obtained, then stores the value 3 to location Z.
- Processor C loads from location Z obtaining the value 3, executes an **msync**, then loads from location X.

#### 6.2.2.2 Lock acquisition and import barriers

An import barrier is an instruction or instruction sequence that prevents memory accesses caused by instructions following the barrier from being performed before memory accesses that acquire a lock have been performed. An import barrier can be used to ensure that a shared data structure protected by a lock is not accessed until the lock has been acquired. An **msync** can always be used as an import barrier, but the approaches shown below generally yield better performance because they order only the relevant memory accesses.

### 6.2.2.2.1 Acquire lock and import shared memory

If **lwarx** and **stwcx.** are used to obtain the lock, an import barrier can be constructed by placing an **isync** immediately following the loop containing the **lwarx** and **stwcx.**. The following example uses the compare and swap primitive (see [Section C.1.1: Synchronization primitives](#)) to acquire the lock.

This example assumes that the address of the lock is in GPR 3, the value indicating that the lock is free is in GPR 4, the value to which the lock should be set is in GPR 5, the old value of the lock is returned in GPR 6, and the address of the shared data structure is in GPR 9.

```
loop: lwarxr6, 0, r3      # load lock and reserve
      cmpw r4, r6        # skip ahead if
      bne- wait         # lock not free
      stwcx. r5, 0, r3   # try to set lock
      bne- loop         # loop if lost reservation
      isync             # import barrier
      lwz r7, data1(r9)  # load shared data
      .
      .
wait: ...                #wait for lock to free
```

The second **bne-** does not complete until CR0 has been set by the **stwcx.**. The **stwcx.** does not set CR0 until it has completed (successfully or unsuccessfully). The lock is acquired when the **stwcx.** completes successfully. Together, the second **bne-** and the subsequent **isync** create an import barrier that prevents the load from `data1` from being performed until the branch is resolved to be not taken.

### 6.2.2.2.2 Obtain pointer and import shared memory

If **lwarx** and **stwcx.** are used to obtain a pointer into a shared data structure, an import barrier is not needed if all the accesses to the shared data structure depend on the value obtained for the pointer. The following example uses the fetch and add primitive (see [Section C.1.1: Synchronization primitives](#)) to obtain and increment the pointer.

In this example, it is assumed that the address of the pointer is in GPR 3, the value to be added to the pointer is in GPR 4, and the old value of the pointer is returned in GPR 5.

```
loop:  lwarx      r5, 0, r3      # load pointer and reserve
      add        r0, r4, r5     # increment the pointer
      stwcx.     r0, 0, r3     # try to store new
value  bne-      loop          # loop if lost reservation
      lwz       r7, data1(r5)   # load shared data
```

The load from `data1` cannot be performed until the **lwarx** loads the pointer value into GPR 5. The load from `data1` may be performed out of order before the **stwcx.**. But if the **stwcx.** fails, the branch is taken and the value returned by the load from `data1` is discarded. If the **stwcx.** succeeds, the value returned by the load from `data1` is valid even if the load is performed out of order, because the load uses the pointer value returned by the instance of the **lwarx** that created the reservation used by the successful **stwcx.**.

An **isync** could be placed between the **bne-** and the subsequent **lwz**, but no **isync** is needed if all accesses to the shared data structure depend on the value returned by the **lwarx**.

### 6.2.2.3 Atomic memory references

The Book E architecture defines the Load Word and Reserve Indexed (**lwarx**) and the store word conditional indexed (**stwcx.**) instructions to provide an atomic update function for a single, aligned word of memory. These instructions can be used to develop a rich set of multiprocessor synchronization primitives. Note that atomic memory references constructed using **lwarx/stwcx.** instructions depend on the presence of a coherent memory system for correct operation. These instructions should not be expected to provide atomic access to noncoherent memory.

The **lwarx** instruction performs a load word from memory operation and creates a reservation for the same reservation granule that contains the accessed word. Reservation granularity is implementation-dependent.

The **lwarx** instruction makes a nonspecific reservation with respect to the executing processor and a specific reservation with respect to other masters. This means that any subsequent **stwcx.** executed by the same processor, regardless of address, cancels the reservation. Also, any bus write or invalidate operation from another processor to an address that matches the reservation address cancels the reservation.

## 6.3 Cache model

A cache model in which there is one cache for instructions and another cache for data is called a 'Harvard-style' cache. This is the model assumed by Book E, for example in the descriptions of the cache management instructions in [Section 4: Instruction model](#). Book E allows the following additional cache models are defined by the EIS:

- Unified cache, in which a cache is shared by both instructions and data
- Multi-level caches, which must support the programming model implied by a Harvard-style cache.

A processor is not required to maintain copies of storage locations in the instruction cache that are consistent with modifications to those storage locations (that is, modifications by store instructions).

In general, a location in the data cache is considered to be modified in that cache if the location has been modified (for example, by a store instruction) and the modified data has not been written to main storage. The only exception to this rule is described in [Section 6.2.1.1: Write-through attribute](#).

Cache management instructions are provided so that programs can manage the caches when needed. For example, program management of the caches is needed when a program generates or modifies code that will be executed (i.e., when the program modifies data in storage and then attempts to execute the modified data as instructions). Cache management instructions are also useful in optimizing the use of memory bandwidth in such applications as graphics and numerically intensive computing. The functions performed by these instructions depend on the storage attributes associated with the specified storage location.

Cache management instructions allow the program to do the following.

- Give a hint that a block of storage should be copied to the instruction cache, so that the copy of the block is more likely to be in the cache when subsequent accesses to the block occur, thereby reducing delays (**icbt**)
- Invalidate the copy of storage in an instruction cache block (**icbi**)

- Discard prefetched instructions (**isync**)
- Invalidate the copy of storage in a data cache block (**dcbi**)
- Give a hint that a block of storage should be copied to the data cache, so that the copy of the block is more likely to be in the cache when subsequent accesses to the block occur, thereby reducing delays (**dcbt**, **dcbtst**)
- Allocate a data cache block and set the contents of that block to zeros, but no-operation if no access is allowed to the data cache block and do not cause any exceptions (**dcba**)
- Set the contents of a data cache block to zeros (**dcbz**)
- Copy the contents of a modified data cache block to main storage (**dcbst**)
- Copy the contents of a modified data cache block to main storage and make the copy of the block in the data cache invalid (**dcbf**).

### 6.3.1 Cache programming model

This section summarizes the register and instructions defined to support the cache model. Full descriptions of these resources are provided in [Section 3: Register model](#), and [Section 4: Instruction model](#).

#### 6.3.1.1 Cache model registers

The EIS cache model implements the following registers and register fields:

- Machine state register (MSR). Defines the processor state (that is, enabling and disabling of interrupts and debugging exceptions, enabling and disabling of address translation for instruction and data memory accesses, enabling and disabling some APUs, and specifying whether the processor is in supervisor or user mode). EIS storage defines the user cache locking enable bit (MSR[UCLE]) as part of the cache line locking APU.

Book E and the EIS define the MSR fields described in [Table 183](#). The MSR is described in detail in [Section 3.6.1: Machine state register \(MSR\)](#).

Table 183. Storage related MSR fields

Bits	Name	Description
37	UCLE	(EIS-defined) User-mode cache lock enable. Used to restrict user-mode cache-line locking by the operating system. 0 Any cache lock instruction executed in user-mode takes a cache-locking exception and data storage interrupt and sets either ESR[DLK] or ESR[ILK]. This allows the operating system to manage and track the locking/unlocking of cache lines by user-mode tasks. 1 Cache-locking instructions can be executed in user-mode and they do not take a DSI for cache-locking (they may still take a DSI for access violations though).
58	IS	(Book E–defined) Instruction address space 0 The processor directs all instruction fetches to address space 0 (TS = 0 in the relevant TLB entry). 1 The processor directs all instruction fetches to address space 1 (TS = 1 in the relevant TLB entry).
59	DS	(Book E–defined) Data address space 0 The processor directs data memory accesses to address space 0 (TS = 0 in the relevant TLB entry). 1 The processor directs data memory accesses to address space 1 (TS = 1 in the relevant TLB entry).

- Exception syndrome register (ESR). The ESR provides a syndrome to differentiate between different kinds of exceptions that can generate the same interrupt type. When such an interrupt is generated, bits corresponding to the exception that generated the interrupt are set and all other ESR bits are cleared. Other interrupt types do not affect ESR contents. The ESR does not need to be cleared by software.
- Book E and the EIS defines the storage-related ESR fields described in [Table 184](#). The ESR is described in detail in [Section 3.9.1.8: Exception syndrome register \(ESR\)](#).

Table 184. Exception syndrome register (ESR) definition

Bits	Name	Syndrome	Interrupt types
39	FP	(Book E–defined) Floating-point operations	Alignment, data storage, data TLB, program
40	ST	(Book E–defined) Store operation	Alignment, data storage, data TLB error
42	DLK	Defined by cache line locking APU. Instruction cache locking attempt. Set when a DSI occurs because a <b>dcbtls</b> , <b>dcbtstls</b> , or <b>dcblc</b> was executed in user mode (MSR[PR] = 1) while MSR[UCLE] = 0. 0 Default 1 DSI occurred on an attempt to lock line in data cache when MSR[UCLE] = 0.	Data storage
43	ILK	Defined by cache line locking APU. Instruction cache locking attempt. Set when a DSI occurs because an <b>icbtls</b> or <b>icblc</b> was executed in user mode (MSR[PR] = 1) while MSR[UCLE] = 0. 0 Default 1 DSI occurred on an attempt to lock line in instruction cache when MSR[UCLE] = 0.	Data storage
44	APU	(Book E–defined) Auxiliary processor operation.	Alignment, data storage, data TLB, program
46	BO	Byte-ordering exception. Defined by Book E and the VLE extension.	Data storage, instruction storage
56	SPE	Defined by SPE, embedded floating-point APU. SPE/embedded floating-point exception bit 0 Default 1 Any exception caused by an SPE/embedded floating-point instruction occurred.	Data storage, Data TLB error, Alignment, SPE unavailable, Embedded FP unavailable, Embedded FP data, Embedded FP round
58	VLEMI	Defined by VLE extension. VLEMI indicates that an interrupt was caused by a VLE instruction. VLEMI is set on an exception associated with execution or attempted execution of a VLE instruction. 0 The instruction page associated with the instruction causing the exception does not have the VLE attribute set or the VLE extension is not implemented. 1 The instruction page associated with the instruction causing the exception has the VLE attribute set and the VLE extension is implemented.	Data storage, Data TLB error, Instruction storage, Program, System Call, Alignment, SPE unavailable, Embedded FP unavailable, Embedded FP data, Embedded FP round

Table 184. Exception syndrome register (ESR) definition (continued)

Bits	Name	Syndrome	Interrupt types
62	MIF	Defined by the VLE extension. MIF indicates that an interrupt was caused by a misaligned instruction fetch ( $NIA_{62} \neq 0$ ) and the VLE attribute is cleared for the page or the second half of a 32-bit VLE instruction caused an instruction TLB error. 0 Default. 1 $NIA_{62} \neq 0$ and the instruction page associated with $NIA$ does not have the VLE attribute set or the second half of a 32-bit VLE instruction caused an instruction TLB error.	Instruction TLB error, Instruction Storage
63	XTE	External transaction error. An external transaction reported an error but the error was handled precisely by the core. SRR0 holds the address of the instruction that initiated the transaction. 0 Default. No external transaction error was precisely detected. 1 An external transaction reported an error that was precisely detected.	Instruction storage, Data storage

- L1 cache control and status registers (L1CSR0–L1CSR1).
  - L1CSR0 provides general control and status for the processor's primary data cache. If a processor implements a unified L1 cache, L1CSR0 applies to the unified cache and L1CSR1 is not implemented. See [Section 3.11.1: L1 cache control and status register 0 \(L1CSR0\)](#).
  - L1CSR1 provides general control and status for the processor's primary instruction cache. If a processor implements a unified L1 cache, L1CSR0 applies to the unified cache and L1CSR1 is not implemented. See [Section 3.11.2: L1 cache control and status register 1 \(L1CSR1\)](#).
- L1 cache configuration registers (L1CFG0)
  - L1CFG0 provides configuration information for the processor's primary data cache. If a processor implements a unified cache, L1CFG0 applies to the unified cache and L1CFG1 is not implemented. See [Section 3.11.3: L1 cache configuration register 0 \(L1CFG0\)](#).
  - L1CFG1 provides configuration information for the processor's primary instruction cache. If a processor implements a unified cache, L1CFG0 applies to the unified cache and L1CFG1 is not implemented. L1CFG1 allows software to identify the organization and capabilities of the primary instruction cache. [Section 3.11.4: L1 cache configuration register 1 \(L1CFG1\)](#).



### 6.3.1.2 Cache model instructions

The Book E PowerPC architecture defines instructions for controlling both the instruction and data caches (when they exist).

- Data Cache Block Touch (**dcbt**)
- Data Cache Block Touch for Store (**dcbtst**)
- Data Cache Block Zero (**dcbz**)
- Data Cache Block Store (**dcbst**)
- Data Cache Block Flush (**dcbf**)
- Data Cache Block Allocate (**dcba**)
- Data Cache Block Invalidate (**dcbi**)
- Instruction Cache Block Invalidate (**icbi**)
- Instruction Synchronize (**isync**)
- Instruction Cache Block Touch (**icbt**)

These instructions are described in [Section 4.3.1.18: User-level cache instructions](#), and [Section 4.3.2.3: Supervisor-level cache instruction](#). Note that the behavior of many of these instructions is determined by the value of the cache target operand (CT). See [Section 6.3.1.3: CT instruction field](#). [Section 6.4.8.4: Permission control and cache management instructions](#), describes conditions in which cache control instructions can generate protection violations.

The cache block locking APU, defined by the EIS, adds the following instructions:

- Data Cache Block Lock Clear (**dcblc**)
- Data Cache Block Touch and Lock Set (**dcbtls**)
- Data Cache Block Touch for Store and Lock Set (**dcbtstls**)
- Instruction Cache Block Lock Clear (**icblc**)
- Instruction Cache Block Touch and Lock Set (**icbtls**)

These instructions are described in [Section 9.1: Cache line locking APU](#).

### 6.3.1.3 CT instruction field

Instructions having a CT (cache target) field for specifying a cache hierarchy use the value 0 to specify the primary cache. ST devices interpret this operand as follows:

- CT = 0 indicates the L1 cache.
- CT = 1 indicates the I/O cache. (Note that some versions of the e500 documentation refer to the I/O cache as a frontside L2 cache.)
- CT = 2 indicates a backside L2 cache.

## 6.3.2 Primary (L1) cache model

This section describes the L1 cache model defined by the EIS.

### 6.3.2.1 Types

Primary caches may separate instruction and data caches into two separate structures (commonly known as Harvard architecture), or they may provide a unified cache combining instructions and data. Caches are physically tagged.

### 6.3.2.2 Storage attributes and coherency

Primary data caches must support the storage attributes defined by Book E with the following advisory:

*Note:* *The primary data cache may be implemented not to snoop (that is, not coherent with transactions outside the processor). System software is then responsible to maintain coherency. Thus the setting of the M attribute is meaningless. The preferred implementation provides snooping for primary data caches.*

Primary instruction caches must support the storage attributes defined by Book E with the following advisory:

- The guarded attribute should be ignored for instruction fetch accesses. To prevent speculative fetch accesses to guarded memory, software should mark those pages as no-execute.
- The cache may be implemented not to snoop (that is, not coherent with transactions outside the processor). System software is then responsible to maintain coherency. The preferred implementation does not provide snooping for primary instruction caches.

As with other memory-related instructions, the effects of cache management instructions on memory are weakly-ordered. If the programmer must ensure that cache or other instructions have been performed with respect to all other processors and system mechanisms, an **msync** must be placed after those instructions.

## 6.4 Storage model

This section describes the storage model as it is defined by Book E and by the EIS.

### 6.4.1 Storage programming model

This section summarizes the register and instructions defined to support the cache model. Full descriptions of these resources are provided in [Chapter 3: Register model](#), and [Chapter 4: Instruction model](#).

#### 6.4.1.1 Storage model registers

This section provides an overview of the registers used for programming the MMU. Full descriptions are provided in [Section 3.12: MMU registers](#). These registers consist of the following:

- Process ID registers (PID0–PID2) are used by system software to identify TLB entries that are used by the processor to accomplish address translation for loads, stores, and instruction fetches. Book E defines one PID register (PID synonymous with PID0). The EIS defines 14 additional PID registers, PID1 through PID14. A implementation may choose to provide any number of PIDs up to a maximum of 15. The number of PIDs implemented is indicated by the value of MMUCFG[NPIDS] and the number of bits implemented in each PID register is indicated by the value of MMUCFG[PIDSIZE]. PID values are used to construct virtual addresses for accessing memory (see [Section 6.4.6: Address translation](#)).
- MMU assist registers (MAS0–MAS7) are used to transfer data to and from the TLB arrays. Software uses **mfspir** and **mtspir** to read and write MAS registers. Executing **tlbre** causes the TLB entry specified by MAS0[TLBSEL,ESEL] and MAS2[EPN] to be copied to the MAS registers. Conversely, execution of a **tlbwe** instruction causes the

TLB entry specified by MAS0[TLBSEL,ESEL] and MAS2[EPN] to be written with the MAS register contents. Hardware can also update MAS registers on the occurrence of an instruction or data TLB error interrupt or as the result of a **tlbsx**.

All MAS registers are supervisor level, and all except MAS5 and MAS7 must be implemented. MAS7 is not required if the processor supports 32 bits or less of physical address. Implementing MAS5 is implementation dependent.

Processors are required to implement only the necessary bits of any multiple-bit MAS register field such that only the resources supplied by the processor are represented. Any non-implemented bits in a field should have no effect when writing and should always read as zero. For example, a processor that implements only two TLB arrays would likely implement only the lower-order MAS0[TLBSEL] bits.

- MAS0, contains fields for identifying and selecting a TLB entry.
- MAS1, contains fields for selecting a TLB entry during translation.
- MAS2, contains fields for specifying the effective page address and the storage attributes for a TLB entry.
- MAS3, contains fields for specifying the real page address and the permission attributes for a TLB entry.
- MAS4, contains fields for specifying default information to be pre-loaded on certain MMU-related exceptions.
- The optional MAS5 register, contains fields for specifying PID values to be used when searching TLB entries with the **tlbsx** instruction.
- MAS6, contains fields for specifying PID and AS values used when the **tlbsx** instruction is used to search TLB entries.

MAS7, contains the high-order address bits of the RPN for implementations that support more than 32 bits of physical address. Implementations that support 32 bits or fewer do not implement MAS7.

- MMU configuration register (MMUCFG), provides configuration information about the MMU.
- TLB configuration registers (TLB<sub>n</sub>CFG). One TLB<sub>n</sub>CFG register, is implemented to provide information about each TLB implemented. TLB0CFG corresponds to TLB0, TLB1CFG corresponds to TLB1, etc.
- MMU control and status register (MMUCSR0), is used for general control of the MMU including flash invalidation of the TLB arrays and page sizes for programmable fixed size arrays. For TLB arrays with programmable fixed sizes, the TLB<sub>n</sub>\_PS fields allow software to specify the page size.

#### 6.4.1.2 Storage model instructions

The address translation mechanism is defined in terms of TLBs and page table entries (PTEs) Book E processors use to locate the logical-to-physical address mapping for a particular access. [Table 103](#) describes the operation of the TLB instructions, which are summarized as follows:

- TLB Invalidate Virtual Address Indexed (**tlbivax**)
- TLB Read Entry (**tlbre**)
- TLB Search Indexed (**tlbsx**)
- TLB Synchronize (**tlbsync**)
- TLB Write Entry (**tlbwe**)

## 6.4.2 The storage architecture

This section describes the storage model as it is defined by Book E and by the ST EIS.

### 6.4.2.1 Book E storage architecture

The memory management approach defined by the Book E EIS is suited for desktop applications and has the simplicity and flexibility necessary for embedded applications. Book E supports demand-paged virtual memory as well as a variety of other management schemes that depend on precise control of effective-to-real address translation and flexible memory protection. Address translation misses and protection faults cause precise exceptions. Sufficient information is available to correct the fault and restart the faulting instruction.

Each program on a 32-bit implementation can access  $2^{32}$  bytes of effective address (EA) space, subject to limitations imposed by the operating system. In a typical Book E system, each program's EA space is a subset of a larger virtual address (VA) space managed by the operating system.

Each effective (logical) address is translated to a real (physical) address before being used to access physical memory or an I/O device. Hardware does this by using the address translation mechanism described in [Section 6.4.6](#). The operating system manages the physically addressed resources of the system by setting up the tables used by the address translation mechanism.

The Book E architecture divides the effective address space into pages. The page represents the granularity of effective address translation, permission control, and memory/cache attributes. Up to 12 page sizes (1, 4, 16, 64, or 256 Kbytes; 1, 4, 16, 64, or 256 Mbytes; or 1 Gbyte) may be simultaneously supported. For an effective-to-real address translation to exist, a valid entry for the page containing the effective address must be in a translation lookaside buffer (TLB). Addresses for which no TLB entry exists cause TLB miss exceptions (instruction or data TLB error interrupts).

The instruction addresses generated by a program and the addresses used by load, store, and cache management instructions are effective addresses. However, in general, the physical memory space may not be large enough to map all the virtual pages used by the currently active applications. With support provided by hardware, the operating system can attempt to use the available real pages to map enough virtual pages for an application. If a sufficient set is maintained, paging activity is minimized, therefore maximizing performance.

The operating system can restrict access to virtual pages by selectively granting permissions for user-state read, write, and execute, and supervisor-state read, write, and execute on a per-page basis. These permissions can be set up for a particular system (for example, program code might be execute-only, data structures may be mapped as read/write/no-execute) and can also be changed by the operating system based on application requests and operating system policies.

### 6.4.2.2 EIS storage architecture

The standard for Book E MMUs establishes a common way of implementing Book E processors to provide a programming model that is consistent across all products in the family. Having a standard reduces the software efforts required in porting to a new processor because the common programming model minimizes implementation differences. Thus, the standard defines configuration information for features such as TLBs, caches, and other entities that have standard forms, but differing attributes (like cache sizes and

associativity) such that a single software implementation can be created that works efficiently for all implementations of a class.

The Book E MMU standard defines functions and structures that are visible to the execution model of the processor. These consist of the following definitions:

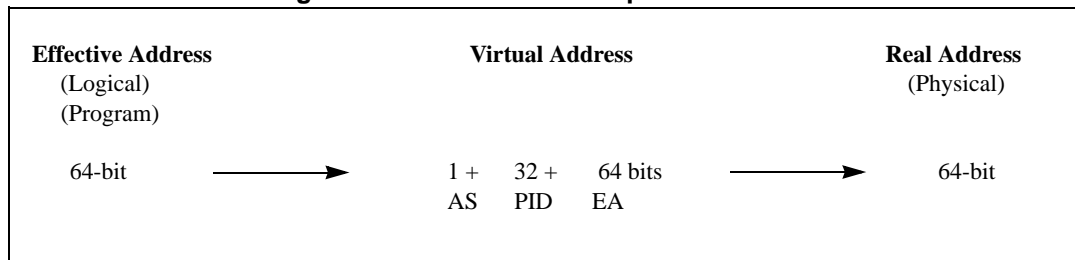
- The TLB, from a programming point of view, consists of zero or more TLB arrays, each of which may have differing characteristics.
- The logical-to-physical address translation mechanism
- Methods and effects of changing and manipulating TLB arrays
- Configuration information available to the operating system that describes the structure and form of the TLB arrays and translation mechanism

To assist or accelerate translation, implementations may contain other TLB structures not visible to the programming model. These structures and the methods for using them are not explicitly defined in the architecture or the ST standard, but they may be considered at the operating system level because they may affect an implementation’s performance.

### 6.4.3 Virtual address (VA)

Book E defines a virtual address space composed of the effective address of an access, the 1-bit current address space (AS) of the access and the 32-bit process ID (PID) of an access, as shown in [Figure 79](#). The following subsections describe the selection of AS and PID for an effective address, both used to construct the virtual address for an access.

**Figure 79. Virtual Address Space in Book E**

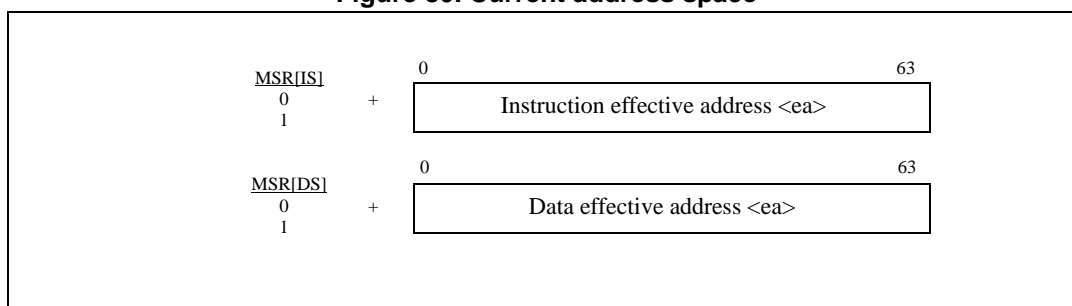


### 6.4.4 Address spaces

Instruction accesses are generated by sequential instruction fetches or due to a change in program flow (branches and interrupts). Data accesses are generated by load, store, and cache management instructions.

The Book E architecture defines two address spaces for instruction accesses and two address spaces for data accesses. The current address space for instruction or data accesses is determined by the value of MSR[IS] and MSR[DS], respectively, as shown in [Figure 80](#).

Figure 80. Current address space



If the type of translation performed is an instruction fetch, the value of the AS bit is taken from the contents of MSR[IS]. If the type of translation performed is a load, store, or other data translation including target addresses of software-initiated instruction fetch hints and locks (**icbt**, **icbtlc**, **icbtlc**) the value of the AS bit is taken from the contents of MSR[DS].

The address space indicator (MSR[IS] or MSR[DS], as appropriate) is used in addition to the effective address generated by the processor for translation into a physical address by the TLB mechanism.

Because MSR[IS] and MSR[DS] are cleared when an interrupt occurs, an address space value of zero can be used to denote interrupt-related address spaces, or possibly all system software address spaces; an address space value of one can be used to denote non-interrupt-related address spaces, or possibly all user address spaces.

Software Note: Although system software is free to use address space bits as it sees fit, on an interrupt, the MSR[IS] and MSR[DS] are cleared. This encourages software to use address space 0 for system software and address space 1 for user software.

#### 6.4.4.1 Instruction address spaces

The two effective instruction address spaces are defined by the value of MSR[IS], and instruction fetch addresses are translated from the effective address space specified by the current value of MSR[IS]. Changing the value of MSR[IS] is considered a context-altering operation, requiring a context synchronization operation to follow it. When a context synchronizing event occurs, any prefetched instructions are discarded and instructions are refetched using the then-current state of MSR[IS] and the then-current program counter. See [Section 4.2.3.6: Context synchronization](#), for more information on the definition of context synchronizing events.

Instructions are not fetched from memory designated by the TLB mechanism as no-execute (UX = 0 or SX = 0). If the effective address of the current instruction is mapped to no-execute memory, an instruction storage interrupt (ISI) is generated.

Note that mapping a page as no-execute does not affect instruction caches in the system (or any instructions resident in unified caches). Thus, if an instruction is loaded into a cache when its effective address is mapped to execute permitted memory, and the execute permissions for that page are later changed to no-execute, any instructions fetched before the no-execute mapping remain in the cache until explicitly evicted by an **icbi** instruction or through the cache's replacement policy. However, attempted execution of such instructions still results in an ISI. Thus, for example, the operating system can change the designation of an application's instruction pages to no-execute without having to first flush instruction cache blocks that map to these pages.

### 6.4.4.2 Data address spaces


The two effective data address spaces are defined by the value of MSR[DS] and data is accessed to/from the effective address space specified by the current value of MSR[DS]. As is the case with MSR[IS], changing the value of MSR[DS] is considered a context-altering operation, requiring a context synchronization operation to follow it. When a context synchronizing event occurs, subsequent accesses are made using the new state of MSR[DS] (see [Section 4.2.3.6: Context synchronization](#)).

Data can be read from a page, provided the user read (UR) permission bit is set in the TLB for a user access, or the supervisor read (SR) bit is set for a supervisor access. Likewise, data write access permissions are determined by the user write (UW) and supervisor write (SW) permission bits. If permissions are violated, the appropriate interrupt is taken.

### 6.4.5 Process ID

As described in [Section 3.12.1: Process ID registers \(PID0–PIDn\)](#), Book E defines that a PID value be associated with each effective address (instruction or data) generated by the processor. At the Book E level, one 32-bit PID register maintains the PID value for the current process. This value is used to construct a virtual address for accessing memory.

Figure 81. Current PID Value



System software uses PIDs to identify TLB entries that the processor uses to translate addresses for loads, stores, and instruction fetches. PID contents are compared to the TID field in TLB entries as part of selecting appropriate TLB entries for address translation. PID values are used to construct virtual addresses for accessing memory. Note that individual processors may not implement all 14 bits of the process ID field.

Book E defines one PID register that holds the PID value for the current process. ST devices may implement from 1 to 15 PID registers. The number of PIDs implemented is indicated by the value of MMUCFG[NPIDS]. Consult the user documentation for the implementation to determine if other PID registers are implemented.

PID registers are more fully described in [Section 3.12.1: Process ID registers \(PID0–PIDn\)](#).

Software Note: The suggested PID usage is for PID0 to denote private mappings for a process and for other PIDs to handle mappings that may be common to multiple processes. This method allows for processes sharing address space to also share TLB entries if the shared address space is mapped at the same virtual address in each process.

#### 6.4.5.1 Process ID (PID) registers

The Book E architecture specifies that a process ID (PID) value be associated with each EA (instruction or data) generated by the processor.

System software uses PIDs to identify TLB entries that the processor uses to translate addresses for loads, stores, and instruction fetches. PID contents are compared to the TID field in TLB entries as part of selecting appropriate TLB entries for address translation. PID

values are used to construct virtual addresses for accessing memory. Note that individual processors may not implement all 14 bits of the process ID field.

Book E defines one PID register that holds the PID value for the current process. ST devices may implement from 1 to 15 PID registers. The number of PIDs implemented is indicated by the value of MMUCFG[NPIDS]. Consult the user documentation for the implementation to determine if other PID registers are implemented.

The 15 PID registers supported by the EIS are implemented as SPR registers set by system software, and collectively reflect the process ID of the currently executing context. The system maintains multiple PID values in order to allow the sharing of TLB entries for pages that are shared among multiple execution contexts. For example, system software may assign PID0 to contain the unique process ID (for private mappings for the current processes) and may assign PID1 to contain the unique process ID for a common set of shared libraries.

Note that Book E defines the value of all zeros for a TID field in a TLB entry as an entry that is globally shared. Thus, when PID values (up to 12 bits for ST devices) are compared to the TID fields in the TLB arrays for matches, if a TLB entry contains all zeros in the TID field, it globally matches all PID values. PID registers are more fully described in [Section 3.12.1: Process ID registers \(PID0–PIDn\)](#).

#### 6.4.5.2 Address space identifiers

The AS bit is the address space identifier. Thus there are two possible address spaces, 0 and 1. The value of the AS bit is determined by the type of translation performed and from the contents of the MSR when an address is translated. If the type of translation performed is an instruction fetch, the value of the AS bit is taken from the contents of MSR[IS]. If the type of translation performed is a load, store, or other data translation including target addresses of software initiated instruction fetch hints and locks (**icbt**, **icbtls**, **icbtlc**) the value of the AS bit is taken from the contents of MSR[DS]. The AS bit is defined by Book E.

*Note: Although system software is free to use address space bits as it sees fit, it should be noted that on interrupt, the MSR[IS] and MSR[DS] bits are cleared. This encourages software to use address space 0 for system software and address space 1 for user software.*

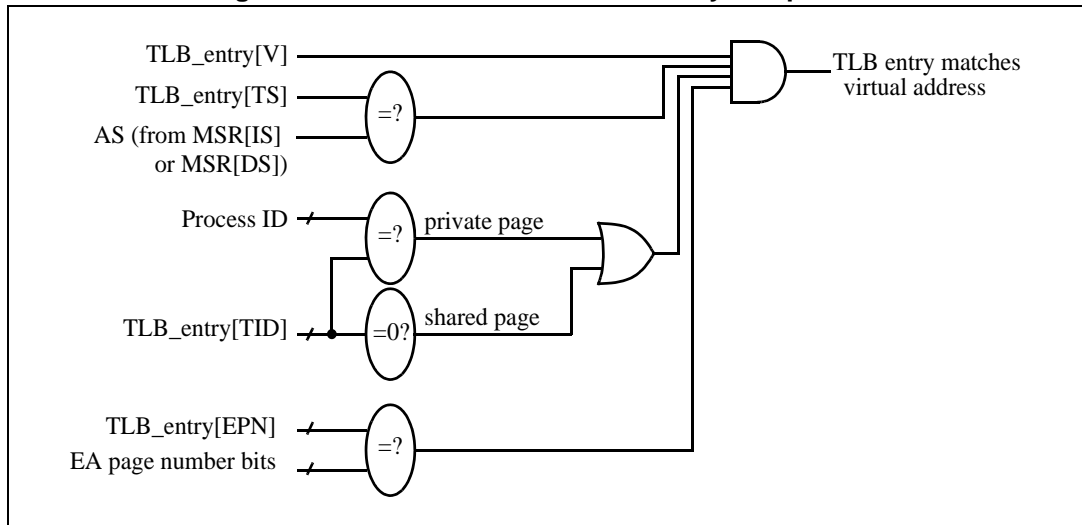
#### 6.4.6 Address translation

The effective address (EA) is the untranslated address for an instruction fetch address or for a data address that is calculated as a result of a load, store, or cache management instruction. The EA, concatenated with the MSR[IS] or MSR[DS] address space (AS) value, is compared to the appropriate number of bits of the EPN field (depending on the page size) and the TS field of the TLB entry. If a match occurs, that TLB entry is a candidate for a translation match. In addition to a match in the EPN field and TS, a matching TLB entry must match with the current process ID of the access.

[Figure 82](#) shows the translation match logic for the effective address plus its attributes (collectively called the virtual address) and how it is compared with the corresponding fields in the TLB entries.

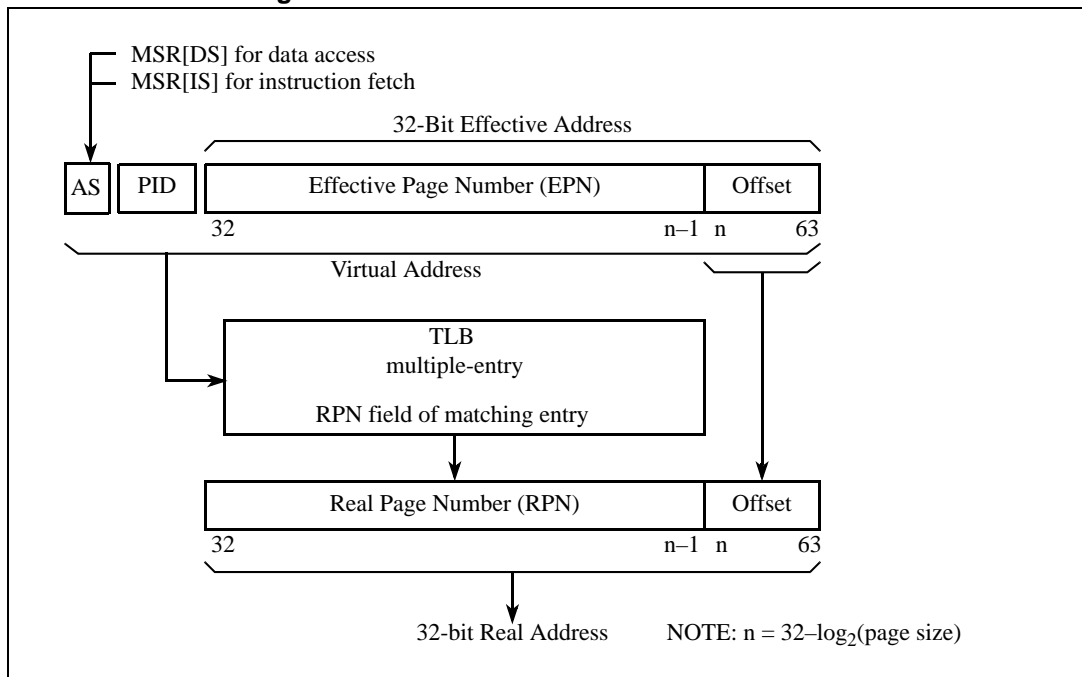


**Figure 82. Virtual address and TLB-entry comparison**



The generation of the physical address occurs as shown in [Figure 83](#).

**Figure 83. Effective-to-real address translation**



The EA combines with the AS and each PID register to form one virtual address for each unique PID register value. Also, an implicit virtual address is formed using a PID value of 0. Thus the following virtual addresses (VAs) are formed:

$$VA0 \leftarrow AS \parallel 0 \parallel EA$$

$$VA1 \leftarrow AS \parallel PID0 \parallel EA$$

...

$$VAN+1 \leftarrow AS \parallel PIDn \parallel EA$$

Note that a PID register containing a 0 value (or the same value as another PID register) forms a non-unique VA. Duplicate VAs are ignored.

Each of the unique VAs are compared to all the valid TLB entries by comparing specific fields of each TLB entry to each of the VAs. The fields of each valid (TLB[V] = 1) TLB entry are combined to form a set of matching TLB address (TAs):

$$TA \leftarrow TLB_{TS} \parallel TLB_{TID} \parallel TLB_{EPN} \parallel 120$$

Each TA is compared to all VAs under a mask based on the page size (TLB[SIZE]) of the TLB entry. The mask of the comparison of the EA and EPN portions of the virtual and translation addresses is computed as follows:

$$mask \leftarrow \sim(1024 \ll (2 * TLB_{SIZE})) - 1$$

where the number of bits in the mask is equal to the number of bits in a TA (or VA). If a TA matches any VA the TLB entry is said to match. If more than one TA/VA match occurs, it is considered a serious programming error and the results are undefined. The recommended behavior is that a machine check interrupt is taken.

Once a match occurs the matching TLB is used for access control, storage attributes, and effective to real address translation. Access control, storage attributes, and address translation are defined by Book E (additional storage attributes are defined within this document).

## 6.4.7 Address translation and the ST EIS

Translating an effective address to a real address is defined by Book E to require four elements:

- The address space value. Depending on the type of translation (instruction or data), MSR[IS] or MSR[DS] is used.
- The TLB entries in the TLB arrays
- The effective address being translated

The following subsections describe these elements as they are further defined by the EIS.

### 6.4.7.1 Match criteria for TLB entries

TLB arrays contain TLB entries that are used to match any address presented for translation. All TLB entries for any given implementation are candidates for any given translation. The TLB itself is unordered with respect to the various elements used in address translations, and regardless of implementation, should be considered to perform the translation comparison with all entries in parallel.

There should be only one valid matching translation for a given effective address, PID value, and address space value. If the TLB contains more than one matching entry, it is considered a programming error, and the behavior of any such translation is undefined. In this case, the processor is likely to enter checkstop state or take a machine check interrupt.

The following fields are compared in the TLB entries:

- V—The matching entry must have the V bit set.
- TS—The address space identifier used for translation. The appropriate bit of MSR[IS] or MSR[DS] must match the TS bit for a matching entry.
- TID—The contents of a PID register must match the TID field of a matching entry, or the TID field must be all zeros for a matching entry.

- EPN—The appropriate number of bits (depending on the page size) of the effective address being translated is compared to the EPN field of the TLB entry.

If a match occurs on all the fields listed above, the physical address is formed by replacing the effective page number in the effective address with the value in the RPN field of the matching TLB entry. The number of bits in the page number depends on the page size for that TLB entry.

#### 6.4.7.2 Translation algorithms

The following algorithm describes how translation operates at the ST Book E level:

```

ea = effective address
if translation is an instruction address then
  as = MSR[IS]
else
  // data address translation
  as = MSR[DS]
for all TLB entries
  if ! TLBV then
    next // compare next TLB entry
  if as != TLBTS then
    next
  if TLBTID == 0 then
    goto pid_match
  for all PID registers
    if this PID register == TLBTID then
      goto pid_match
  endfor
  next // no PIDs matched

pid_match: // translation match
  mask = ~(1024 << (2 * TLBTSIZE)) - 01
  if (ea & mask) != TLBEPN then
    next // no address match
  real address = TLBRPN | (ea & ~mask) // real address computed
end translation -- success
endfor
end translation -- tlbmiss

```

The algorithm for the granting of permission is as follows:

```

if MSRPR == 0 then
  x = TLBSX
  r = TLBSR
  w = TLBSW
else
  x = TLBUX
  r = TLBUR
  w = TLBUW

if instruction fetch address then
  if x == 0 then
    Instruction Storage Interrupt
  else // data access
    if data read (load) then

```

```

    if r == 0 then
        Data Storage Interrupt
    else
        // write access (store)
        if w == 0 then
            Data Storage Interrupt

```

### 6.4.7.3 Access control

If address translation results in a match (hit), the matching TLB entry is used to perform access control (permission checks). These checks are based on the privilege level of the access (MSR[PR]) and the type of access (fetch for execute, read for loads, and write for stores). The TLB entry's permission bits (TLB[US,SX,UW,SW,UR,SR]) determine if the operation should succeed. If permission is denied, execution of the instruction is suppressed and an instruction storage interrupt or data storage interrupt occurs as defined in Book E. Software uses the ESR, SRR0, and the DEAR to determine the type of operation attempted and then must perform a TLB search if updating the TLB is desired.

The algorithm for determining access control is as follows:

```

if MSRPR = 0 then
    x ← TLBSX
    r ← TLBSR
    w ← TLBSW
else
    x ← TLBUX
    r ← TLBUR
    w ← TLBUW

if instruction_fetch & x = 0 then
    take instruction storage interrupt
else if load & r = 0 then
    take data storage interrupt
else if store & w = 0 then
    take data storage interrupt
else
    access permitted

```

### 6.4.7.4 Physical (real) address generation

If permission checking is successful, the real address is formed by combining the TLB[RPN] with the lower order offset bits of the EA based on the page size of the TLB entry.

$$\text{mask} \leftarrow \sim(1024 \ll (2 * \text{TLB}_{\text{SIZE}})) - 1$$

$$\text{real\_address} \leftarrow ((\text{TLB}_{\text{RPN}} \ll 12) \& \text{mask}) | (\text{EA} \& \sim\text{mask})$$

Where mask contains the same number of bits as a real address. The real address is then used to access the memory subsystem using the TLB[ACM,VLE,W,I,M,G,E] fields from the TLB entry to determine how the location should be accessed.

### 6.4.7.5 Page size and effective address bits compared

The page size defined for a TLB entry determines how many bits of the effective address are compared with the corresponding EPN field in the TLB entry as shown in [Table 185](#).

Table 185. Page size and EPN field comparison

SIZE Field	Page Size ( $4^{\text{SIZE}}$ Kbytes)	EA to EPN Comparison (Bits 32–53; $2 \times \text{SIZE}$ )
0b0000	1 Kbyte	EA[32–53] = ? EPN[32–53]
0b0001	4 Kbyte	EA[32–51] = ? EPN[32–51]
0b0010	16 Kbyte	EA[32–49] = ? EPN[0–49]
0b0011	64 Kbyte	EA[32–47] = ? EPN[32–47]
0b0100	256 Kbyte	EA[32–45] = ? EPN[32–45]
0b0101	1 Mbyte	EA[32–43] = ? EPN[32–43]
0b0110	4 Mbyte	EA[32–41] = ? EPN[32–41]
0b0111	16 Mbyte	EA[32–39] = ? EPN[32–39]
0b1000	64 Mbyte	EA[32–37] = ? EPN[32–37]
0b1001	256 Mbyte	EA[32–35] = ? EPN[32–35]
0b1010	1 Gbyte	EA[32–33] = ? EPN[32–33]

#### 6.4.7.6 Permission attribute comparison

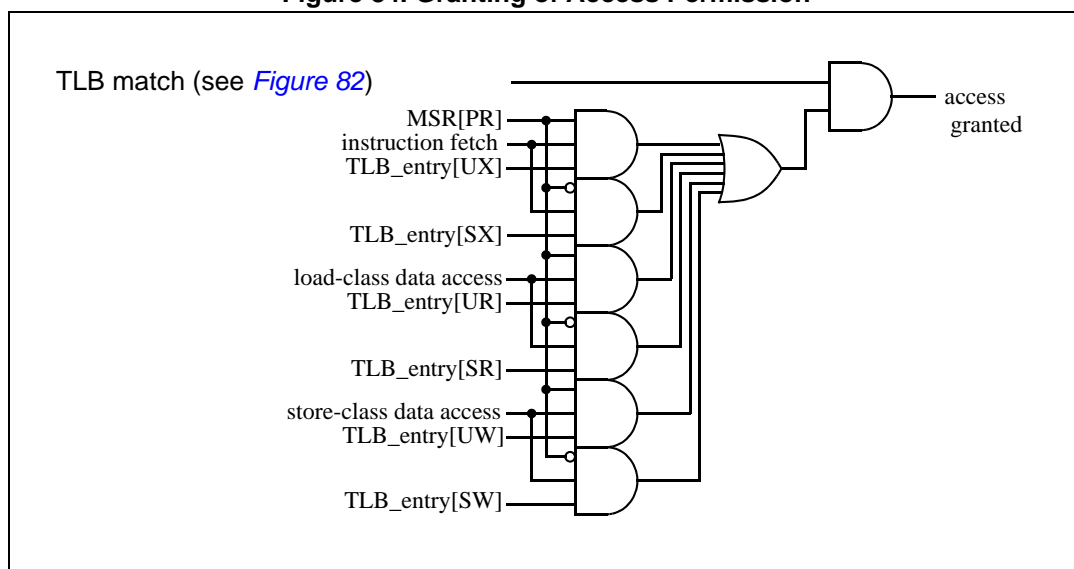
As part of the translation process, the selected TLB entry provides the access permission bits (UX, SX, UW, SW, UR, SR), and memory/cache attributes (U0, U1, U2, U3, W, I, M, G, and E) for the access. These bits specify whether or not the access is allowed and how the access is to be performed.

If a matching TLB entry has been identified, Book E provides an access permission mechanism that selectively grants shared access, grants execute access, grants read access, grants write access, and prohibits access to areas of memory based on a number of criteria. Book E defines the permission bits in TLB entries as follows:

- SR—Supervisor read permission
- SW—Supervisor write permission
- SX—Supervisor execute permission
- UR—User read permission
- UW—User write permission
- UX—User execute permission

If the virtual address translation comparison with TLB entries was successful, the permission bits for the matching entry are checked as shown in [Figure 84](#). If the access is not allowed by the access permission mechanism, the processor generates an instruction or data storage interrupt (ISI or DSI).

Figure 84. Granting of Access Permission



The permission attributes defined by Book E are defined in detail in [Section 6.4.8](#).

### 6.4.7.7 Page size and real address generation

If no virtual address match occurs, the translation fails and a TLB miss exception occurs. Depending on the access type (instruction or data address), either the instruction TLB error interrupt or the data TLB error interrupt is taken.

Otherwise, the real page number (RPN) field of the matching TLB entry provides the translation for the effective address of the access. Based on the setting of the SIZE field of the matching TLB entry, the RPN field replaces the corresponding most-significant n bits of the effective address where  $n = 32 - \log_2(\text{page size})$ . Note that the untranslated bits must be zero in the RPN field.

Table 186. Real address generation

Size field	Page size ( $4^{\text{SIZE}}$ Kbytes)	RPN bits required to be equal to 0	Real address
0b0000	1 Kbyte	none	RPN[32–53]    EA[54–63]
0b0001	4 Kbyte	RPN[52–53] = 0	RPN[32–51]    EA[52–63]
0b0010	16 Kbyte	RPN[50–53] = 0	RPN[32–49]    EA[50–63]
0b0011	64 Kbyte	RPN[48–53] = 0	RPN[32–47]    EA[48–63]
0b0100	256 Kbyte	RPN[46–53] = 0	RPN[32–45]    EA[46–63]
0b0101	1 Mbyte	RPN[44–53] = 0	RPN[32–43]    EA[44–63]
0b0110	4 Mbyte	RPN[42–53] = 0	RPN[32–41]    EA[42–63]
0b0111	16 Mbyte	RPN[40–53] = 0	RPN[32–39]    EA[40–63]
0b1000	64 Mbyte	RPN[38–53] = 0	RPN[32–37]    EA[38–63]
0b1001	256 Mbyte	RPN[36–53] = 0	RPN[32–35]    EA[36–63]
0b1010	1 Gbyte	RPN[34–53] = 0	RPN[32–33]    EA[34–63]

## 6.4.8 Permission attributes

The permission attributes defined in Book E are shown in [Table 187](#) and described in the following subsections.

**Table 187. Permission control for instruction, data read, and data write accesses**

Access type	MSR[PR]	TLB[UX]		TLB[SX]		TLB[UR]		TLB[SR]		TLB[UW]		TLB[SW]	
		0	1	0	1	0	1	0	1	0	1	0	1
Instruction fetch	0	—	—	ISI	√	—	—	—	—	—	—	—	—
	1	ISI	√	—	—	—	—	—	—	—	—	—	—
Data read (load)	0	—	—	—	—	—	—	DSI	√	—	—	—	—
	1	—	—	—	—	DSI	√	—	—	—	—	—	—
Data write (store)	0	—	—	—	—	—	—	—	—	—	—	DSI	√
	1	—	—	—	—	—	—	—	—	DSI	√	—	—

### 6.4.8.1 Execute access permission

The UX and SX bits of the TLB entry control execute access to the corresponding page.

Instructions may be fetched and executed from a page in memory if MSR[PR] = 1 (user mode) if the UX access control bit for that page is set. If the UX access control bit is cleared, instructions from that page are not fetched and they are not placed into any cache while the processor is in user mode.

Instructions may be fetched and executed from a page in memory if MSR[PR] = 0 (supervisor mode) and the SX access control bit for that page is set. If the SX access control bit is cleared, instructions from that page are not fetched and are not placed into any cache while the processor is in supervisor mode.

If the sequential execution model calls for the execution of an instruction from a page that is not enabled for execution (that is, UX = 0 when MSR[PR] = 1 or SX = 0 when MSR[PR] = 0), an execute access control exception-type instruction storage interrupt (ISI) is taken.

### 6.4.8.2 Read access permission

The UR and SR bits of the TLB entry control read access to the corresponding page.

Load operations (including load-class cache management instructions) are permitted from a page in memory while the processor is in user mode (MSR[PR] = 1) if the UR access control bit for that page is set. If the UR access control bit is cleared, execution of the load instruction is suppressed and a read access control exception-type data storage interrupt (DSI) is taken.

Similarly, load operations (including load-class cache management instructions) are permitted from a page in memory if MSR[PR] = 0 (supervisor mode) and the SR access control bit for that page is set. If the SR access control bit is cleared, execution of the load instruction is suppressed and a read access control exception-type data storage interrupt (DSI) is taken.

### 6.4.8.3 Write access permission

The UW and SW bits of the TLB entry control write access to the corresponding page.

Store operations (including store-class cache management instructions) are permitted to a page in memory if MSR[PR] = 1 (user mode) and the UW access control bit for that page is set. If the UW access control bit is cleared, execution of the store instruction is suppressed and a write access control exception-type data storage interrupt (DSI) is taken.

Similarly, store operations (including store-class cache management instructions) are permitted to a page in memory if MSR[PR] = 0 (supervisor mode) and the SW access control bit for that page is set. If the SW access control bit is cleared, execution of the store instruction is suppressed and a write access control exception-type data storage interrupt (DSI) is taken.

### 6.4.8.4 Permission control and cache management instructions

The **dcbi** and **dcbz** instructions are treated as stores because they can change data (or cause loss of data by invalidating a modified line). As such, they both can cause write access control exception-type DSIs.

The **dcba** instruction is treated as a store because it can also change data. As such, it can also cause a write access control exception. However, these exceptions do not result in a data storage interrupt and if a permission violation occurs, the instruction execution completes, but the allocate operation is merely cancelled (essentially, a no-op).

The **icbi** instruction is treated as a load with respect to permissions checking. As such, it can cause a read access control exception-type data storage interrupt.

The **dcbt**, **dcbtst**, and **icbt** instructions are treated as loads with respect to permissions checking. As such, they can cause read access control exceptions. However, such exceptions do not result in data storage interrupts and if a permission violation occurs, the instruction execution completes, but the operation is cancelled (essentially, a no-op).

The **dcbf** and **dcbst** instructions are treated as loads with respect to permissions checking. Flushing or storing a line from the cache is not considered a store because the store has already been performed to update the cache and the **dcbf** or **dcbst** instruction is only updating the copy in main memory. Like load instructions, these instructions can cause read access control exception-type data storage interrupts.

[Table 188](#) summarizes exception cases caused by the cache management instructions due to permissions violations.

**Table 188. Permission control and cache instructions**

Instruction	Can cause read permission violation exception?	Can cause write permission violation exception?
<b>dcba</b>	No	Yes <sup>(1)</sup>
<b>dcbf</b>	Yes	No
<b>dcbi</b>	No	Yes
<b>dcbst</b>	Yes	No
<b>dcbt</b>	Yes <sup>(1)</sup>	No
<b>dcbtst</b>	Yes <sup>(1)</sup>	No
<b>dcbz</b>	No	Yes



**Table 188. Permission control and cache instructions (continued)**

Instruction	Can cause read permission violation exception?	Can cause write permission violation exception?
<b>icbi</b>	Yes	No
<b>icbt</b>	Yes <sup>(1)</sup>	No

1. **dcba**, **dcbt**, **dcbtst**, and **icbt** may cause a read access control exception but does not result in a data storage interrupt (DSI).

### 6.4.8.5 Permissions control and string instructions

When the string length is zero, neither **lswx** nor **stswx** can cause data storage interrupts due to permissions violations.

### 6.4.8.6 Use of permissions to maintain page history

The Book E architecture TLB entry definition does not define bits for maintaining page history information. The U0–U3 bits in the TLB entries can be used by software for storing history information, but implementations may ignore these bits internally.

Page changed bit status can be implemented in the system software by disabling write permissions to all pages. The first attempt to write to the page results in a data storage interrupt. At this point, system software can record the page changed bit in memory, update the TLB entry permission to allow writes to that page, and return to the user program allowing further writes to the page to proceed without exception.

### 6.4.8.7 Crossing page boundaries

Care must be taken with single instruction accesses (load/stores) that cross page boundaries. Examples are **lmw** and **stmw** instructions and misaligned accesses on implementations that support misaligned load/stores. Architecturally, each of the parts of the access that cross the natural boundary of the access size (half word, word, double word) are treated separately with respect to exception conditions. Additionally, these types of instructions may optionally partially complete. For example, a store word instruction that crosses a page boundary because it is misaligned to the last half word of a page might actually store the first 16 bits because the access was permitted, but produce a DSI or data TLB error exception because the second 16 bits in the next page were not valid or they were protected. An implementation may choose to suppress the first 16-bit store or perform it.

### 6.4.9 Translation lookaside buffer (TLB) arrays

The MMU contains up to four TLB arrays, which are on-chip storage areas for holding TLB entries. A TLB entry contains effective-to-physical address mappings for loads, stores, and instruction fetches. A TLB array must contain TLB entries that share the same characteristics and contains zero or more TLB entries. Each TLB entry has specific fields that can be addressed by the corresponding fields in the MMU assist registers (see [Section 3.12.5: MMU assist registers \(MAS0–MAS7\)](#)). Each implemented TLB array has an associated configuration register (TLBnCFG) describing the size and attributes of the TLB entries in that array. See [Section 3.12.4: TLB configuration registers \(TLBnCFG\)](#).

The architected fields of a TLB entry are described in [Table 189](#).

Table 189. TLB entry

Name	Description
V	Valid bit. A 1-bit entry that specifies whether this TLB entry is valid for translation. (1 = valid).
TID	Translation ID. Identifies which process ID (PID) that this TLB entry is valid for. Translation IDs are compared with Process IDs (PIDs) during translation to identify which TLB entry to use for translating an address. A TID value of 0 is considered global and matches all PID values.
TS	Translation space. Identifies which address space that this TLB entry is valid for. The translation space field is compared with MSR[IS] for instruction accesses and the MSR[DS] bit for data accesses. This allows for an efficient change of address space when a transition from user mode to supervisor mode occurs. This is a 1 bit field.
SIZE	Page size. Describes the size of the page. An implementation is not required to support variable page sizes or any particular page size. If a TLB array does not support variable size pages (that is, TLBnCFG[AVAIL] = 0) then this field is ignored. Page sizes range from 4 Kbytes to 1 Tbyte in powers of 4. EIS does not support 1-Kbyte page sizes defined in Book E. Page size encoding is defined by Book E.
EPN	Effective page number. Describes the logical or effective starting address of the page. The number of bits that are valid (used in translation) depends on the size of the page and if the processor is a 32- or 64-bit implementation. This field is used to compare with the EA being translated to identify which TLB entry to use for translation. For 32-bit implementations, this field is 2 to 20 bits, depending on the page size (SIZE).
RPN	Real page number. Describes the physical starting address of the page. The real page number is substituted for the effective page number from the address being translated which results in the real address. This field is 2 to 52 bits depending on the page size and the number of bits of real address supported by the implementation.
WIMGE	Storage attributes. Describe the characteristics of any memory/fetch accesses to the page and the subsequent treatment of those data items with respect to the memory subsystem (caches and bus transactions). The WIMGE bits are defined by Book E.
ACM	Alternate coherency mode. Optional. An implementation may optionally support additional coherency models. If such coherency models are provided, they are encoded in this field. ACM values are implementation dependent. The Alternate Coherence Mode is used only when the M bit (from WIMGE) is set.
VLE	Variable length encoding. Optional. If an implementation supports the VLE extension, clearing VLE causes instruction access to this page to decode and execute as PowerPC Book E (and EIS APUs) instructions, and setting VLE causes instruction access to this page to decode and execute as VLE (and EIS APU) instructions.
SR,SW,SX, UR,UW,UX	Permissions. User and supervisor read, write, and execute permission bits. Supervisor and user permission bits are defined by Book E.
U0,U1,U2,U3	User bits. Implementation dependent. Consult the user's manual for any implementation usage. It is strongly recommended to leave these as storage associated with a TLB entry to be used by system software.
IPROT	Invalidation protection. Invalidation protection. This entry is protected from all TLB invalidation mechanisms except the explicit writing of a 0 to the V bit.

## 6.4.10 TLB management

TLB entries are managed by software using the set of MAS registers, which are used to move data between software and the TLB entries, to identify TLB entries, and to provide default values when translation or protection faults occur. See [Section 3.12.5: MMU assist registers \(MAS0–MAS7\)](#).

### 6.4.10.1 TLB configuration information

Information about the configuration for a given TLB implementation is available to system software by reading the contents of the MMU configuration SPRs. These SPRs describe the architectural version of the MMU, the number of TLB arrays, and the characteristics of each TLB array. MMU architecture version number 1 is defined as these SPRs with the field definitions as described in this section.

- MMU configuration register (MMUCFG), implemented by all ST Book E processors, contains basic information about the MMU architecture for each device. TLB configuration registers (TLB $n$ CFG). Implemented by all ST Book E processors for each of the TLBs specified in MMUCFG[NTLBS]. They contain configuration information about each particular TLB. See [Section 3.12.3: MMU configuration register \(MMUCFG\)](#).
- The TLB $n$ CFG number assignment is the same as the value in MAS0[TLBSEL]. For example, TLB0CFG provides configuration information about TLB0, and TLB1CFG provides configuration information about TLB1. See [Section 3.12.4: TLB configuration registers \(TLB \$n\$ CFG\)](#).

### 6.4.10.2 TLB entries

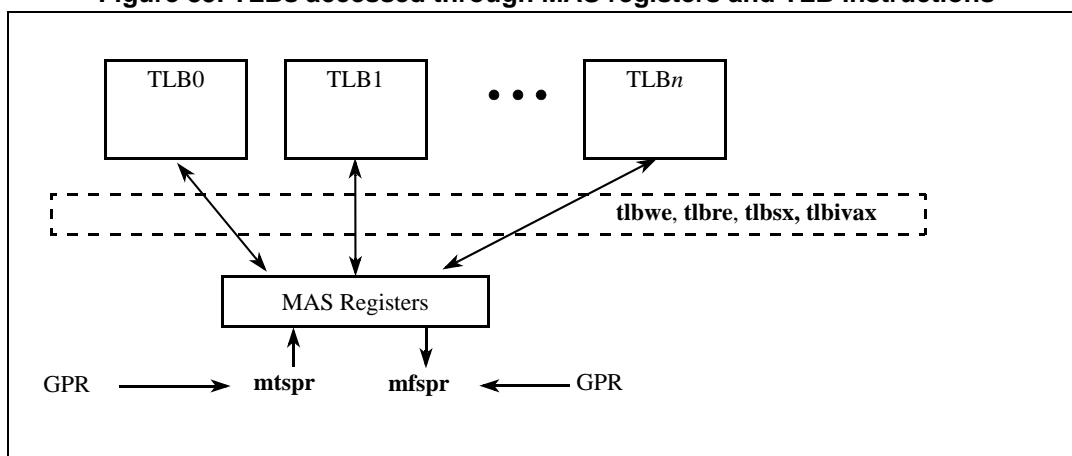
The software-visible TLB is subdivided into zero or more TLB arrays. Each array must contain TLB entries that share the same characteristics. Each TLB array contains one or more TLB entries. Each entry has specific fields that correspond to fields in the seven MMU assist (MAS) registers, described in [Section 3.12.5: MMU assist registers \(MAS0–MAS7\)](#). Some TLB fields are architected in Book E and others are architected in the EIS. Note that Book E architected fields may have restrictions or enhancements imposed by the EIS for the Book E implementations.

The IPROT TLB entry, architected by the EIS, designates TLB entries as protected from certain kinds of invalidation. TLB invalidation and the IPROT field are described further in [Section 6.4.10.6: Invalidating TLB entries](#).

### 6.4.10.3 Reading and writing TLB entries

All TLB entries are updated by executing **tlbwe** instructions. At the time of **tlbwe** execution, the MMU assist registers (MAS0–MAS6), a set of SPRs defined by the EIS, are used to index a specific TLB entry. The MAS registers also contain the information that is written to the indexed entry, such that they serve as the ports into the TLBs, as shown in [Figure 85](#). The contents of the MAS registers are described in [Section 3.12.5: MMU assist registers \(MAS0–MAS7\)](#).

Figure 85. TLBs accessed through MAS registers and TLB instructions



Similarly, TLB entries are read by executing **tlbre** instructions. At the time of **tlbre** execution, the MAS registers are used to index a specific TLB entry and upon completion of the **tlbre** instruction, the MAS registers contain the contents of the indexed TLB entry. To read or write TLB entries, the MAS registers are first loaded by system software using **mtspr** instructions and then the desired **tlbre** or **tlbwe** instructions must be executed.

Note that RA = 0 is a preferred form for **tlbxs** and that some Book E implementations take an illegal instruction exception program interrupt if RA ≠ 0.

#### 6.4.10.4 Reading TLB entries

TLB entries are read by executing **tlbre** instructions. At the time of **tlbre** execution, the MAS registers are used to index a specific TLB entry and upon completion of the **tlbre**, the MAS registers contain the contents of the indexed TLB entry.

Selection of the TLB entry to read is performed by setting MAS0[TLBSEL], MAS0[ESEL] and MAS2[EPN] to indicate the entry to read. MAS0[TLBSEL] selects which TLB the entry should be read from (0 to 3) and MAS2[EPN] selects the set of entries from which MAS0[ESEL] selects an entry. For fully associative TLBs, MAS2[EPN] is not required since the value in MAS0[ESEL] fully identifies the TLB entry. Valid values for MAS0[ESEL] are from 0 to associativity - 1.

The selected TLB entry is then used to update the following fields of the MAS registers: V, IPROT, TID, TS, TSIZE, EPN, ACM, VLE, WIMGE, RPN, U0—U3, & permissions. If the TLB array supports NV, it is used to update the NV field in the MAS registers, otherwise the contents of NV field are undefined. The update of MAS registers as a result of a **tlbre** instruction is summarized in [Table 191](#).

No operands are given for the **tlbre** instruction and the Book E defined implementation dependent field should be treated as a reserved field.

Specifying invalid values for MAS0[TLBSEL] and MAS0[ESEL] produce boundedly undefined results.

#### 6.4.10.5 Writing TLB entries

TLB entries are written by executing **tlbwe** instructions. At the time of **tlbwe** execution, the MAS registers are used to index a specific TLB entry and contain the contents to be written to the indexed TLB entry. Upon completion of the **tlbwe** instruction, the TLB entry contents of the MAS registers are written to the indexed TLB entry.

Selection of the TLB entry to write is performed by setting MAS0[TLBSEL], MAS0[ESEL] and MAS2[EPN] to indicate the entry to write. MAS0[TLBSEL] selects which TLB the entry should be written from (0 to 3) and MAS2[EPN] selects the set of entries from which MAS0[ESEL] selects an entry. For fully associative TLBs, MAS2[EPN] is not used to identify a TLB entry since the value in MAS0[ESEL] fully identifies the TLB entry. Valid values for MAS0[ESEL] are from 0 to associativity minus 1.

The selected TLB entry is then written with following fields of the MAS registers: V, IPROT, TID, TS, TSIZE, EPN, ACM, VLE, WIMGE, RPN, U0—U3, and permissions. If the TLB array supports NV, it is written with the NV value.

The effects of updating the TLB entry are not guaranteed to be visible to the programming model until the completion of a context synchronizing operation. Writing a TLB entry that is used by the programming model prior to a context synchronizing operation produces undefined behavior.

No operands are given for the **tlbwe** instruction and the Book E defined implementation dependent field should be treated as a reserved field.

Specifying invalid values for MAS0[TLBSEL] and MAS0[ESEL] produce boundedly undefined results.

*Note:* Writing TLB entries should be followed by an **isync** or an **rfi** before the new entries are to be used by the programming model.

#### 6.4.10.6 Invalidating TLB entries

TLB entries may be invalidated by any of the following methods:

- A TLB entry can be invalidated as the result of a **tlbwe** instruction that clears MAS0[V] in the entry.
- As a result of a **tlbivax** instruction or from a received broadcast invalidation resulting from a **tlbivax** on another processor in an SMP system.
- As a result of a flash invalidate.

In both multiprocessor and uniprocessor systems, invalidations can occur on a wider set of TLB entries than intended. This are called generous invalidations That is, a virtual address presented for invalidation may invalidate not only the targeted TLB, but also may invalidate other TLB entries, depending on the implementation. This is because parts of the translation mechanism may not be fully specified to the hardware at invalidate time. This is especially true in SMP systems where the invalidation address must be broadcast globally to all processors in the system. Hardware may impose other limitations.

The architecture ensures that the intended TLB is invalidated, but does not guarantee that it is the only one. A TLB entry invalidated by clearing the V bit of the TLB entry by use of a **tlbwe** is guaranteed to invalidate only the addressed TLB entry. However, invalidates occurring from **tlbivax** instructions or from the multiprocessor broadcasts as a result of **tlbivax** instructions may cause generous invalidates.

The architecture provides a method to protect against generous invalidations. This is important, because certain virtual memory regions (most notably, the code memory region that serves as the exception handler for MMU faults) must be properly mapped to for forward progress to occur. If this region does not have a valid mapping, an MMU exception cannot be handled because the first address of the interrupt handler causes another MMU exception. To prevent this, the architecture specifies an IPROT bit for TLB entries. Setting the MAS0[PROT] protects the corresponding TLB entry from invalidations resulting from **tlbivax** instructions, as a result of broadcast invalidation from another processor in an SMP

system, or from flash invalidations. TLB entries with the IPROT field set can be invalidated only by explicitly writing the TLB entry and specifying a 0 for MAS1[V].

*Note:* **Software Note:** Not all TLB arrays in a given implementation implement the IPROT attribute. It is likely that implementations that are suitable for demand page environments implement it for only a single array, while not implementing it for other arrays.

**Software Note:** Operating systems must use great care when using protected (IPROT) TLB entries, particularly in SMP systems. An SMP system that contains TLB entries on other processors requires a cross-processor interrupt or some other synchronization mechanism to assure that each processor performs the required invalidation by writing its own TLB entries.

### 6.4.10.7 Invalidations using tlbivax:


The **tlbivax** instruction provides a virtual address as a target for invalidation. EA[0–51] are used to find a TLB entry with a matching EPN field. The page size of the TLB entry is used to mask the low order bits in the comparison. The comparison is performed only for TLB entries in the specified TLB array, do not have the IPROT attribute set (if supported by the TLB array), and are valid. The AS bit does not participate in the comparison. The EA specified by the **rA** and **rB** operands in the **tlbivax** instruction contains fields in the lower order bits to augment the invalidation to specific TLB arrays and to flash invalidate those arrays. Note that TLB entry invalidations resulting from **tlbivax** instructions do not invalidate any entry that has IPROT = 1 unless the specified TLB array does not support the IPROT attribute. The encoding of the EA used by **tlbivax** is shown in [Table 190](#).

*Note:* **Software Note:** To ensure a TLB entry that is not protected by IPROT is invalidated if software does not know which TLB array the entry is in, software should issue a **tlbivax** instruction targeting each TLB in the implementation with the EA to be invalidated.

**Software Note:** The preferred form of the **tlbivax** instruction contains the entire EA in **rB** and zero in **rA**. Some implementations may take an Unimplemented Instruction exception if **rA** is non-zero.

EA format for **tlbivax** is shown below.

#### EA Format for tlbivax



[Table 190](#) describes EA fields for **tlbivax**.

**Table 190. Fields for EA format of tlbivax**

Field	Name	Comments or function when set
0–51	EA <sub>0:51</sub>	The upper bits of the address to invalidate.
52–58	—	Reserved, should be cleared.
59–60	TLB	Selects TLB array for invalidation. 00 TLB0 01 TLB1 10 TLB2 11 TLB3

Table 190. Fields for EA format of `tlbivax`

Field	Name	Comments or function when set
61	IA	Invalidate all entries in selected TLB array.
62–63	—	Reserved, should be cleared.

Flash invalidations using MMUCSR0:

All entries in a TLB array may be flash invalidated using the MMUCSR0 register. Flash invalidation of an array is started when the corresponding flash invalidate bit is set in MMUCSR0 (MMUCSR0[TLBn\_FI]). The flash invalidation is complete when the corresponding flash invalidate bit is cleared by the processor. Writing a 0 to a flash invalidate bit in MMUCSR0 has no effect. Note that TLB entry invalidations resulting from MMUCSR0 flash invalidations do not invalidate any entry that has IPROT = 1 unless the specified TLB array does not support the IPROT attribute.

#### 6.4.10.8 Searching TLB Entries

Software may search the MMU by using the `tlbsx` instruction that is provided by Book E. The `tlbsx` instruction uses PID values and an AS value from the MAS registers instead of the PID registers and the MSR. This allows software to search address spaces that differ from the current address space defined by the PID registers. This is useful for TLB fault handling.

To properly execute a search for a TLB, software loads MAS5 and MAS6 registers with PID and AS values to search for. These are MAS6[SPID0], MAS6[SPID1], MAS5[SPID2], MAS5[SPID3], MAS6[SAS]. Software then executes a `tlbsx` instruction. The search performs the same TA to VA comparison described in [Section 6.4.6](#), except that the PID and AS values are taken from the MAS registers. If a matching, valid TLB entry is found, the MAS register are loaded with the information from that TLB entry as if the TLB entry was read from a `tlbre` instruction. Software can examine the MAS1[V] bit to determine if the search was successful. Successful searches cause the valid bit to be set. [Table 191](#) summarizes the update of MAS registers as a result of a `tlbsx` instruction.

The preferred form of the `tlbsx` is `rA = 0`. Some implementations may take an unimplemented instruction exception or an illegal instruction exception if `rA != 0`.

#### 6.4.10.9 TLB replacement hardware assist

The architecture provides mechanisms to accelerate software in creating and updating TLB entries when MMU related exceptions occur. This is called TLB replacement hardware assist. Hardware updates the MAS registers on the occurrence of a data TLB error interrupt or instruction TLB error interrupt.

When a TLB error exception (miss) occurs, MAS0, MAS1, and MAS2 are automatically updated using the defaults specified in MAS4 as well as the AS and EPN values corresponding to the access that caused the exception. MAS6 is updated to set MAS6[SPID0] to the value of PID0 and MAS6[SAS] to the value of MSR[DS] or MSR[IS] depending on the type of access that caused the TLB error. In addition, if MAS4[TLBSELD] identifies a TLB array that supports NV (next victim), MAS0[ESEL] is loaded with a value that hardware believes represents the best TLB entry to victimize to create a new TLB entry and MAS0[NV] is updated with the TLB entry index of what hardware believes to be the next victim. Thus MAS0[ESEL] identifies the current TLB entry to be replaced, and MAS0[NV] points to the next victim. When software writes the TLB entry, MAS0[NV] is written to the

TLB array. The algorithm used by the hardware to determine which TLB entry should be targeted for replacement is implementation dependent.

The automatic update of MAS registers sets up all the necessary fields for creating a new TLB entry with the exception of RPN, the U0–U3 attribute bits, and the permission bits. With the exception of the upper 32 bits of RPN and the page attributes (should software desire to specify changes from the default attributes), all remaining fields are located in MAS3, requiring only the single MAS register manipulation by software before writing the TLB entry.

For ISI and DSI related exceptions, the MAS registers are not updated. Software must explicitly search the TLB to find the appropriate entry.

The update of MAS registers through TLB replacement hardware assist is summarized in [Table 191](#).

**Table 191. MAS register update summary**

MAS field updated	Value loaded on event			
	TLB error interrupt	tlbsx hit	tlbsx miss	tlbre
MAS0[TLBSEL]	MAS4[TLBSELD]	TLB array that hit	MAS4[TLBSELD]	—
MAS0[ESEL]	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> hardware hint, <b>else</b> undefined	Number of entry that hit	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> hardware hint, <b>else</b> undefined	—
MAS0[NV]	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> next hardware hint, <b>else</b> undefined	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> hardware hint, <b>else</b> undefined	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> next hardware hint, <b>else</b> undefined	<i>if</i> MAS4[TLBSELD] supports next victim <b>then</b> hardware hint, <b>else</b> undefined
MAS1[V]	1	1	0	TLB[V]
MAS1[I][PROT]	0	TLB[I][PROT]	0	TLB[I][PROT]
MAS1[TID]	<i>if</i> PID[MAS4[TIDSELD]] implemented <b>then</b> PID[MAS4[TIDSELD]] <b>else</b> 0	TLB[TID]	MAS6[SPID0]	TLB[TID]
MAS1[TS]	MSR[IS] or MSR[DS]	TLB[TS]	MAS6[SAS]	TLB[TS]
MAS1[TSIZE]	MAS4[TSIZED]	TLB[SIZE]	MAS4[TSIZED]	TLB[SIZE]
MAS2[EPN]	EA <sub>0:51</sub>	TLB[EPN]	—	TLB[EPN]
MAS2[ACM]	MAS4[ACMD]	TLB[ACM]	MAS4[ACMD]	TLB[ACM]
MAS2[VLE]	MAS4[VLED]	TLB[VLE]	MAS4[VLED]	TLB[VLE]
MAS2[W]	MAS4[WD]	TLB[W]	MAS4[WD]	TLB[W]
MAS2[I]	MAS4[ID]	TLB[I]	MAS4[ID]	TLB[I]
MAS2[M]	MAS4[MD]	TLB[M]	MAS4[MD]	TLB[M]
MAS2[G]	MAS4[GD]	TLB[G]	MAS4[GD]	TLB[G]
MAS2[E]	MAS4[ED]	TLB[E]	MAS4[ED]	TLB[E]



Table 191. MAS register update summary (continued)

MAS field updated	Value loaded on event			
	TLB error interrupt	tlbsx hit	tlbsx miss	tlbre
MAS3[RPN]	0	TLB[RPN] (bits 32:51)	0	TLB[RPN] (bits 32:51)
MAS3[U0,U1,U2,U3]	—	TLB[U0,U1,U2,U3]	—	TLB[U0,U1,U2,U3]
MAS3[UX,SX,UW,SW,UR,SR]	0	TLB[UX,SX,UW,SW,UR,SR]	0	TLB[UX,SX,UW,SW,UR,SR]
MAS4	—	—	—	—
MAS5	—	—	—	—
MAS6[SPID0]	PID0	—	—	—
MAS6[SPID1]	—	—	—	—
MAS6[SAS]	MSR[IS] or MSR[DS]	—	—	—
MAS7[RPN]	0	TLB[RPN] (bits 0–31)	0	TLB[RPN] (bits 0–31)

### 6.4.11 MAS registers and exception handling

When translation-related exceptions occur, hardware preloads the MAS registers with information that the interrupt handler likely needs to handle the fault. For a TLB miss exception, some MAS register fields are loaded with default information specified in MAS4. System software should set up the default information in MAS4 before allowing exceptions. In most cases, system software sets this up once depending on its scheme for handling page faults. This simplifies translation-related exception handling. The following subsections detail specific MAS register fields and the contents loaded for each exception type.

#### 6.4.11.1 TLB miss exception types

The Book E architecture defines that a TLB miss exception is caused when a virtual address for an access does not match with that of any on-chip TLB entry. This condition causes one of the following:

- An instruction TLB error interrupt
- A data TLB error interrupt

##### 6.4.11.1.1 Instruction TLB error interrupt settings

An instruction TLB error interrupt occurs when the virtual address associated with an instruction address (fetch) does not match any valid entry in the TLB (that is, the address for the instruction cannot be translated). In addition to the values automatically written to the MAS registers (described in [TLB miss exception MAS register settings](#)), SRR0 contains the address of the instruction that caused the instruction TLB error. This SRR0 value is used to identify the EA for handling the exception as well as the address to return to when system software has resolved the exception condition by writing a new TLB entry.

### 6.4.11.1.2 Data TLB error interrupt settings

A data TLB error interrupt occurs when the virtual address associated with a data reference from a load, store, or cache management instruction does not match any valid entry in the TLB (that is, the address of the data item of a load or store instruction cannot be translated). In addition to the values automatically written to the MAS registers (described in [TLB miss exception MAS register settings](#)), the effective address of the data access that caused the exception is automatically loaded in the data exception address register (DEAR). Also, SRR0 contains the address of the instruction that caused the data TLB error and its value is used to identify the address to return to when system software has resolved the exception condition (by writing a new TLB entry).

#### TLB miss exception MAS register settings

When either an instruction or data TLB error interrupt occurs, the TLB information and selection fields of the MAS registers are loaded with default values from other MAS registers to assist in processing the exception. The intention is that the common case of a page fault generally requires only system software to load the RPN (corresponding to the physical address that will be used for this page), and the access permissions and the defaults can be used for the remaining MAS fields.

The processor may use the next victim (NV) field from the TLB array to select which TLB entry should be used for the new translation. The method used to select the candidate TLB for replacement (the next victim) is implementation-dependent and may vary on different Book E implementations. In any case, software is free to choose any TLB entry for the replacement (software can overwrite the value in MAS0[ESEL]).

The EIS defines the fields set in the MAS registers at exception time for an instruction or data TLB error interrupt as shown in [Table 192](#).

**Table 192. MAS settings for an instruction or data TLB error interrupt**

MAS Field	Value
TLBSEL	Set to value in TLBSELD (default). This defines the TLB array to be used for the new TLB entry that will be written.
ESEL	May be set to an implementation-dependent value, usually based on the NV field of the array selected by TLBSELD, if that TLB supports the NV function. If the selected TLB does not support NV, the value loaded into ESEL is undefined.
NV	Set to an implementation-dependent value to select which TLB entry to replace on the next TLB miss. The NV field of the TLB array is updated by the value of NV in the MAS registers when <b>tibwe</b> is executed.
V	Set
IPROT	Cleared
TID	Set to the contents of the PID register referenced by TIDSELD. That is, if TIDSELD contains the value 1, the contents of PID1 are written to the TID field.
TS	Set to the value of the IS or DS bit in the MSR at the time of the exception (that is, the MSR that described the context that was running when the exception occurred).
TSIZE	Set to TSIZED
EPN	Set to the effective page number of the instruction or data address causing the exception. The number of bits of the page number is implementation-dependent, but should be consistent with the TLB array selected if the TLB has a fixed page size. If the TLB array selected by TLBSELD contains variable-sized pages, the value for EPN is undefined.

Table 192. MAS settings for an instruction or data TLB error interrupt

MAS Field	Value
WIMGE and X bits	Set to corresponding default values in the MAS registers
RPN	Cleared
Permissions	SR, UR, UW, SW, UX, SX cleared to 0 (no permissions); note that U0–U3 are unchanged.

All other MAS register values are unchanged.

### 6.4.11.2 Permissions violation exception types

The Book E architecture also defines that a permissions violation exception is caused when an effective address for an access matches with a TLB entry but the permission attributes in the matching TLB entry do not allow the access to proceed, as described in [Section 6.4.8](#). This condition causes an instruction storage interrupt (ISI) or a data storage interrupt (DSI)

#### Instruction storage interrupt settings

An instruction storage interrupt occurs when the effective address associated with an instruction address (fetch) matches a valid entry in the TLB but one of the permission bits in the TLB does not allow the instruction fetch. In addition to the values automatically written to the MAS registers (described in [Permissions violation mas register settings](#)), SRR0 contains the address of the instruction that caused the instruction TLB error and this value is used to identify the effective address for handling the exception as well as the address to return to when system software has resolved the exception condition (by writing a new TLB entry).

#### 6.4.11.2.1 Data storage interrupt settings

A data storage interrupt occurs when the EA associated with a data reference from a load or store instruction matches with a valid entry in the TLB but one of the permission bits in the TLB does not allow the data access. In addition to the values automatically written to the MAS registers (described in [Permissions violation mas register settings](#)), the effective address of the data access that caused the exception is contained in the data exception address register (DEAR). This address is used by system software to identify the address that caused the exception. Also, SRR0 contains the address of the instruction that caused the data storage interrupt and its value is used to identify the address to return to when system software has resolved the exception condition by writing a new TLB entry.

#### Permissions violation mas register settings

When either an instruction or data storage interrupt occurs, only the SPIDx and the SAS fields are automatically loaded into the MAS registers to assist in processing the interrupt. System software is required to then execute a **tlbsx** instruction to load the MAS registers with the TLB entry associated with the instruction address. System software may then make any desired changes to the TLB entry prior to writing it. [Table 193](#) describes the fields set in the MAS registers at exception time for instruction or data storage interrupts.

**Table 193. MAS settings for permissions violation ISI or DSI**

Field	Setting
SPID0	Set to PID0
SPID1	Set to PID1
SPID2	Set to PID2
SPID3	Set to PID3
SAS	Set to the value of MSR[IS] (for an instruction storage interrupt) or MSR[DS] (for a data storage interrupt) at the time of the exception (that is, the MSR that described the context that was running when the exception occurred).

All other MAS register values are unchanged.

**6.4.11.3 MAS register updates for exceptions, tlbsx, and tlbre**

*Table 194* summarizes MAS register fields updates from the perspective of the EIS as a result of various events. Note that the implementations further define how certain MAS fields are set on exceptions.

**Table 194. MMU assist register field updates—EIS definition**

MAS <sub>n</sub> bit/field	Value loaded for each case						
	ITLB/DTLB error	tlbsx hit	tlbsx miss	ISI	DSI	tlbre	tlbwe
TLBSEL	TLBSELD	Which TLB hit	TLBSELD	—	—	—	—
ESEL	If TLBSELD supports NV: TLB0[NV] else, undefined	Number of entry that hit	If TLBSELD supports NV: TLB0[NV] else, undefined	—	—	—	—
NV	If TLBSELD supports NV: next NV (array) else, undefined	If TLBSEL supports NV: NV (array) else, undefined	If TLBSELD supports NV: next NV (array) else, undefined	—	—	If TLBSEL supports NV: NV (array) else, undefined	—
V	1	1	0	—	—	V (array)	—
IPROT	0	If TLB that hits supports IPROT: matched IPROT value; else, 0	0	—	—	If TLB that hits supports IPROT: matched IPROT value; else, 0	—
TID	PID <sub>n</sub> values selected by TIDSELD	TID (array)	SPID0	—	—	TID(array)	—
TS	MSR[IS/DS]	SAS	SAS	—	—	TS (array)	—
TSIZE[0–3]	TSIZED	TSIZE (array)	TSIZED	—	—	TSIZE (array)	—



**Table 194. MMU assist register field updates—EIS definition (continued)**

MASn bit/field	Value loaded for each case						
	ITLB/DTLB error	tbsx hit	tbsx miss	ISI	DSI	tlbre	tlbwe
EPN[32–51]	If TLBSELD has fixed page size: EPN of access else, undefined	EPN (array)	—	—	—	EPN (array)	—
X0, X1 WIMGE	X0D, X1D WIMGED	X0, X1 (array) WIMGE (array)	X0D, X1D WIMGED	—	—	X0, X1 (array) WIMGE (array)	—
RPN[32–51]	Zeros	RPN(array)	Zeros	—	—	RPN (array)	—
PERMIS	Zeros	PERMIS (array)	Zeros	—	—	PERMIS (array)	—
TLBSELD	—	—	—	—	—	—	—
TIDSELD	—	—	—	—	—	—	—
TSIZED	—	—	—	—	—	—	—
WIMGED	—	—	—	—	—	—	—

## 7 Instruction set

This chapter describes the following instructions:

- Book E instructions defined for 32-bit implementations. This includes instructions not implemented in all Book E devices.
- Instructions defined by the EIS, except for the instructions defined by the VLE extension. Full descriptions of these instructions are provided in [Section 14: VLE instruction set](#).

### 7.1 Notation

The following definitions and notation are used throughout this chapter in the instruction descriptions.

**Table 195. Notation conventions**

Symbol	Meaning
$X_p$	Bit p of register/field X
$X_{\text{field}}$	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
$X_{p:q}$	Bits p through q of register/field X
$X_{p\ q\ \dots}$	Bits p, q, ... of register/field X
$\neg X$	The one's complement of the contents of X
Field i	Bits $4 \times i$ through $4 \times i + 3$ of a register
.	As the last character of an instruction mnemonic, this character indicates that the instruction records status information in certain fields of the condition register as a side effect of execution, as described in <a href="#">Section 3.5.1: Condition register (CR)</a> .
	Describes the concatenation of two values. For example, 010    111 is the same as 010111.
$x^n$	x raised to the $n^{\text{th}}$ power
${}^n x$	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 0b0_0000. ${}^n 1$ means a field of n bits with each bit equal to 1. Thus ${}^5 1$ is equivalent to 0b1_1111.
/, //, ///,	A reserved field in an instruction or in a register. Each bit and field in instructions, in status and control registers (such as the XER or FPSCR), and in SPRs is either defined, allocated, or reserved, as described in <a href="#">Section 4.2.1: Classes of instructions</a> .

## 7.2 Instruction fields

[Table 196](#) describes instruction fields.

**Table 196. Instruction field descriptions**

Field	Description
AA (30)	<p>Absolute address bit.</p> <p>0 The immediate field represents an address relative to the current instruction address.            For I-form branch instructions the effective address of the branch target is the value <math>^{32}0 \parallel (CIA+EXTS(LI  0b00))_{32-63}</math>.            For B-form branch instructions the effective address of the branch target is the value <math>^{32}0 \parallel (CIA+EXTS(BD  0b00))_{32-63}</math>.            For I-form branch extended instructions the effective address of the branch target is the value <math>CIA+EXTS(LI  0b00)</math>.            For B-form branch extended instructions the effective address of the branch target is the value <math>CIA+EXTS(BD  0b00)</math>.</p> <p>1 The immediate field represents an absolute address.            For I-form branch instructions the effective address of the branch target is the value <math>^{32}0 \parallel EXTS(LI  0b00)_{32-63}</math>.            For B-form branch instructions the effective address of the branch target is the value <math>^{32}0 \parallel EXTS(BD  0b00)_{32-63}</math>.            For I-form branch extended instructions the effective address of the branch target is the value <math>EXTS(LI  0b00)</math>.            For B-form branch extended instructions the effective address of the branch target is the value <math>EXTS(BD  0b00)</math>.</p>
crbA (11–15)	Used to specify a condition register bit to be used as a source
crbB (16–20)	Used to specify a condition register bit to be used as a source
crD (6–8)	Used to specify a CR or FPSCR field to be used as a target
crS (11–13)	Used to specify a CR or FPSCR field to be used as a source
BI (11–15)	Used to specify a condition register bit to be used as the condition of a branch conditional instruction
BO (6–10)	Used to specify options for branch conditional instructions. See <a href="#">Section 4.3.1.11: Branch and flow control instructions</a> .
crbD (6–10)	Used to specify a CR or FPSCR bit to be used as a target
CT (6–10)	Used by cache touch instructions ( <b>dcbt</b> , <b>dcbtst</b> , and <b>icbt</b> ) to specify the target portion of the cache facility to place the prefetched data or instructions and is implementation-dependent
D (16–31)	Immediate field used to specify a 16-bit signed two's complement integer that is sign-extended to 64 bits
DCRN(16–20  11–15)	Used to specify a device control register for the <b>mtdcr</b> and <b>mfdcr</b> instructions
E (15)	Immediate field used to specify a 1-bit value used by <b>wrtteei</b> to place in MSR[EE] (external input enable bit)
FM (7–14)	Field mask used to identify FPSCR fields that are to be updated by the <b>mtfsf</b> instruction
frA (11–15)	Used to specify an FPR to be used as a source
frB (16–20)	Used to specify an FPR to be used as a source

Table 196. Instruction field descriptions (continued)

Field	Description
frC (21–25)	Used to specify an FPR to be used as a source
frS (6–10)	Used to specify an FPR to be used as a source
frD (6–10)	Used to specify an FPR to be used as a target
CRM (12–19)	Field mask used to identify the condition register fields to be updated by the <b>mtcrf</b> instruction
LI (6–29)	Immediate field specifying a 24-bit signed two's complement integer that is concatenated on the right with 0b00 and sign-extended to 64 bits
LK (31)	LINK bit. Indicates whether the link register (LR) is set. 0Do not set the LR. 1Set the LR. The sum of the value 4 and the address of the branch instruction is placed into the LR.
MB (21–25) and ME (26–30)	Fields used in M-form rotate instructions to specify a 64-bit mask consisting of 1 bits from bit MB+32 through bit ME+32 inclusive and 0 bits elsewhere.
MO (6–10)	Used to specify the subset of memory accesses ordered by a Memory Barrier instruction ( <b>mbar</b> ).
NB (16–20)	Used to specify the number of bytes to move in an immediate Move Assist instruction
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
Rc (31)	Record bit. 0Do not alter the condition register. 1Set condition register field 0 or field 1.
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
SH (16–20)	Used to specify a shift amount in rotate word immediate and shift word immediate instructions
SIMM (16–31)	Immediate field used to specify a 16-bit signed integer
SPRN (16–20  11–15)	Used to specify an SPR for <b>mtspr</b> and <b>mfspir</b> instructions
TO (6–10)	Used to specify the conditions on which to trap. The encoding is described in <a href="#">Table 92: Trap instructions</a> .
U (16–19)	Immediate field used as the data to be placed into a field in the FPSCR
UIMM (16–31)	Immediate field used to specify a 16-bit unsigned integer
XO (21–29, 21–30, 22–30, 26–30, 27–29, 27–30, 28–31)	Extended opcode field

### 7.3 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 195](#) and [Table 196](#) and the RTL-specific conventions in [Table 197](#). See the example in [Figure 86](#). 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 **stwcx.** 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 197. 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)
$+$	Two's complement addition
$-$	Two's complement subtraction, unary minus
$\times$	Multiplication
$\div$	Division (yielding quotient)
$+_{dp}$	Floating-point addition, double precision
$-_{dp}$	Floating-point subtraction, double precision
$\times_{dp}$	Floating-point multiplication, double precision
$\div_{dp}$	Floating-point division quotient, double precision
$+_{sp}$	Floating-point addition, single precision
$-_{sp}$	Floating-point subtraction, single precision
$\times_{sf}$	Signed fractional multiplication. Result of multiplying two quantities having bit lengths $x$ and $y$ taking the least significant $x+y-1$ bits of the product and concatenating a 0 to the least significant bit forming a signed fractional result of $x+y$ bits.
$\times_{si}$	Signed integer multiplication
$\times_{sp}$	Floating-point multiplication, single precision
$\div_{sp}$	Floating-point division, single precision
$\times_{fp}$	Floating-point multiplication to infinite precision (no rounding)
$\times_{ui}$	Unsigned integer multiplication
FPSquareRoot-Double( $x$ )	Floating-point $\sqrt{x}$ , result rounded to double-precision
FPSquareRoot-Single( $x$ )	Floating-point $\sqrt{x}$ , result rounded to single-precision
FPReciprocal-Estimate( $x$ )	Floating-point estimate of $\frac{1}{x}$

Table 197. RTL notation (continued)

Notation	Meaning
FPReciprocal-SquareRoot-Estimate(x)	Floating-point estimate of $\frac{1}{\sqrt{x}}$
Allocate-DataCache-Block(x)	If the block containing the byte addressed by x does not exist in the data cache, allocate a block in the data cache and set the contents of the block to 0.
Flush-DataCache-Block(x)	If the block containing the byte addressed by x exists in the data cache and is dirty, the block is written to main memory and is removed from the data cache.
Invalidate-DataCache-Block(x)	If the block containing the byte addressed by x exists in the data cache, the block is removed from the data cache.
Store-DataCache-Block(x)	If the block containing the byte addressed by x exists the data cache and is dirty, the block is written to main memory but may remain in the data cache.
Prefetch-DataCache-Block(x,y)	If the block containing the byte addressed by x does not exist in the portion of the data cache specified by y, the block in memory is copied into the data cache.
Prefetch-ForStore-DataCache-Block(x,y)	If the block containing the byte addressed by x does not exist in the portion of the data cache specified by y, the block in memory is copied into the data cache and made exclusive to the processor executing the instruction.
ZeroDataCache-Block(x)	The contents of the block containing the byte addressed by x in the data cache is cleared.
Invalidate-Instruction-CacheBlock(x)	If the block containing the byte addressed by x is in the instruction cache, the block is removed from the instruction cache.
Prefetch-Instruction-CacheBlock(x,y)	If the block containing the byte addressed by x does not exist in the portion of the instruction cache specified by y, the block in memory is copied into the instruction cache.
=, ≠	Equals, Not Equals relations
<, ≤, >, ≥	Signed comparison relations
< <sub>u</sub> , > <sub>u</sub>	Unsigned comparison relations
?	Unordered comparison relation
&,	AND, OR logical operators
⊕, ≡	Exclusive OR, Equivalence logical operators ((a≡b) = (a⊕-b))
>>, <<	Shift right or left logical
ABS(x)	Absolute value of x
APID(x)	Returns an implementation-dependent information on the presence and status of the auxiliary processing extensions specified by x
CEIL(x)	Least integer ≥ x
DCREG(x)	Device control register x
DOUBLE(x)	Result of converting x from floating-point single format to floating-point double format
EXTS(x)	Result of extending x on the left with signed bits
EXTZ(x)	Result of extending x on the left with zeros
FPR(x)	Floating-point register x
GPR(x)	General-purpose register x

Table 197. RTL notation (continued)

Notation	Meaning
MASK(x, y)	Mask having 1s in bit positions x through y (wrapping if $x > y$ ) and 0s elsewhere
MEM(x,1)	Contents of the byte of memory located at address x
MEM(x,y) (for $y=\{2,4,8\}$ )	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 of the value being accessed.
MOD(x,y)	Modulo y of x (remainder of x divided by y)
ROTL32(x, y)	Result of rotating the value x left y positions, where x is 32 bits long
SINGLE(x)	Result of converting x from floating-point double format to floating-point single format
SPREG(x)	Special-purpose register x
TRAP	Invoke a trap-type program interrupt
characterization	Reference to the setting of status bits in a standard way that is explained in the text
undefined	An undefined value. The value may vary between implementations and between different executions on the same implementation.
CIA	Current instruction address, the address of the instruction being described in RTL. Used by relative branches to set the next instruction address (NIA) and by branch instructions with LK=1 to set the LR. CIA does not correspond to any architected register.
NIA	Next instruction address, the address of the next instruction to be executed. For a successful branch, the next instruction address is the branch target address: in RTL, this is indicated by assigning a value to NIA. For other instructions that cause non-sequential instruction fetching, the RTL is similar. For instructions that do not branch, and do not otherwise cause instruction fetching to be non-sequential, the next instruction address is CIA+4. NIA does not correspond to any architected register.
if ... then ... else ...	Conditional execution, indenting shows range; else is optional
do	Do loop, indenting shows range. 'To' and/or 'by' clauses specify incrementing an iteration variable, and a 'while' clause gives termination conditions.
leave	Leave innermost do loop, or do loop described in leave statement.

Precedence rules for RTL operators are summarized in [Table 198](#). 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  $a-b-c = (a-b)-c$ .) 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 198. Operator precedence**

Operators	Associativity
Subscript, function evaluation	Left to right
Pre-superscript (replication), post-superscript (exponentiation)	Right to left
unary -, ¬	Right to left
×, ÷	Left to right
+, -	Left to right
	Left to right
=, ≠, <, ≤, >, ≥, <u>, >u, ?	Left to right
&, ⊕, ≡	Left to right
	Left to right
: (range)	None
←	None

### 7.3.1 SPE APU 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.

#### 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
    
```

#### Bit reverse

```

BITREVERSE(value)
result ← 0
mask ← 1
shift ← 31
cnt ← 32
while cnt > 0 then do
  t ← data & mask
  if shift >= 0 then
    result ← (t << shift) | result
  else
    result ← (t >> -shift) | result
  cnt ← cnt - 1
  shift ← shift - 2
  mask ← mask << 1
return result
    
```

## 7.3.2 Embedded floating-point conversion models

The embedded floating-point instructions defined by the signal processing engine (SPE) APU and the single-precision floating-point (SPFP) APUs 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 here as a group of functions that is called from the individual instruction pseudo RTL descriptions.

**Table 199. Conversion models**

Function	Name	Reference
<b>Common functions</b>		
Round a 32-bit value	Round32(fp,guard,sticky)	<a href="#">on page 343</a>
Round a 64-bit value	Round64(fp,guard,sticky)	<a href="#">on page 342</a>
Signal floating-point error	SignalFPError	<a href="#">on page 342</a>
Is a 32-bit value a NaN or Infinity?	Isa32NaNorInfinity(fp)	<a href="#">on page 341</a>
<b>Floating-point conversions</b>		
Convert from single-precision floating-point to integer word with saturation	CnvtFP32ToI32Sat(fp,signed,upper_lower,round,fractional)	<a href="#">on page 343</a>
Convert from double-precision floating-point to integer word with saturation	CnvtFP64ToI32Sat(fp,signed,round,fractional)	<a href="#">on page 345</a>
Convert from double-precision floating-point to integer double word with saturation	CnvtFP64ToI64Sat(fp,signed,round)	<a href="#">on page 346</a>
Convert to single-precision floating-point from integer word with saturation	CnvtI32ToFP32Sat(v,signed,upper_lower,fractional)	<a href="#">on page 348</a>
Convert to double-precision floating-point from integer double word with saturation	CnvtI64ToFP64Sat(v,signed)	<a href="#">on page 349</a>
<b>Integer Saturate</b>		
Integer saturate	SATURATE(ovf,carry,neg_sat,pos_sat,value)	<a href="#">on page 342</a>

### Common embedded floating-point functions

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

```

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) & (fpfrac ≠ 0))
Isa32Infinity(fp)
return ((fpexp = 255) & (fpfrac = 0))

```

```

// Determine if fp value is denormalized
Isa32Denorm(fp)
return ((fpexp = 0) & (fpfrac ≠ 0))
// Determine if fp value is a NaN or Infinity
Isa64NaNorInfinity(fp)
return (fpexp = 2047)
Isa64NaN(fp)
return ((fpexp = 2047) & (fpfrac ≠ 0))
Isa64Infinity(fp)
return ((fpexp = 2047) & (fpfrac = 0))
// Determine if fp value is denormalized
Isa64Denorm(fp)
return ((fpexp = 0) & (fpfrac ≠ 0))

Signal Floating-Point Error
// Signal a Floating-Point Error in the SPEFSCR
SignalFPError(upper_lower, bits)
if (upper_lower = UPPER) then
    bits ← bits << 15
    SPEFSCR ← SPEFSCR | bits
bits ← (FG | FX)
if (upper_lower = UPPER) then
    bits ← bits << 15
    SPEFSCR ← SPEFSCR & ~bits

Round a 32-Bit Value
// Round a result
Round32(fp, guard, sticky)
FP32format fp;
if (SPEFSCRFINXE = 0) then
    if (SPEFSCRFRMC = 0b00) then // nearest
        if (guard) then
            if (sticky | fpfrac[22]) then
                v[0:23] ← fpfrac + 1
                if v[0] then
                    if (fpexp >= 254) then
                        // overflow
                        fp ← fpsign || 0b11111110 || 231
                    else
                        fpexp ← fpexp + 1
                        fpfrac ← v[1:23]
                else
                    fpfrac ← v[1:23]
            else if ((SPEFSCRFRMC & 0b10) = 0b10) then // infinity modes
                // implementation dependent
return fp

Round a 64-Bit Value
// Round a result
Round64(fp, guard, sticky)

```

```

FP32format fp;
if (SPEFSCRFINXE = 0) then
  if (SPEFSCRFRMC = 0b00) then // nearest
    if (guard) then
      if (sticky | fpfrac[51]) then
        v[0:52] ← fpfrac + 1
        if v[0] then
          if (fpexp >= 2046) then
            // overflow
            fp ← fpsign || 0b11111111110 || 521
          else
            fpexp ← fpexp + 1
            fpfrac ← v[1:52]
        else
          fpfrac ← v[1:52]
      else if ((SPEFSCRFRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
    return fp

```

### 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 (fpsign = 1) then
      return 0x80000000
    else
      return 0x7fffffff
  else
    if (fpsign = 1) then
      return 0x00000000
    else
      return 0xffffffff

if (Isa32Denorm(fp)) then
  SignalFPError(upper_lower, FINV)
  return 0x00000000 // regardless of sign

if ((signed = UNSIGN) & (fpsign = 1)) then
  SignalFPError(upper_lower, FOVF) // overflow
  return 0x00000000

```

```

if ((fpexp = 0) & (fpfrac = 0)) then
    return 0x00000000 // all zero values

if (fractional = I) then // convert to integer
    max_exp ← 158
    shift ← 158 - fpexp
    if (signed = SIGN) then
        if ((fpexp ≠ 158) | (fpfrac ≠ 0) | (fpsign ≠ 1)) then
            max_exp ← max_exp - 1
    else // fractional conversion
        max_exp ← 126
        shift ← 126 - fpexp
        if (signed = SIGN) then
            shift ← shift + 1

if (fpexp > max_exp) then
    SignalFPEError(upper_lower, FOVF) // overflow
    if (signed = SIGN) then
        if (fpsign = 1) then
            return 0x80000000
        else
            return 0x7fffffff
    else
        return 0xffffffff

result ← 0b1 || fpfrac || 0b00000000 // add U to frac
guard ← 0
sticky ← 0

for (n ← 0; n < shift; n ← n + 1) do
    sticky ← sticky | guard
    guard ← result & 0x00000001
    result ← result > 1

// Report sticky and guard bits
if (upper_lower = UPPER) then
    SPEFSCRFGH ← guard
    SPEFSCRFXH ← sticky
else
    SPEFSCRFG ← guard
    SPEFSCRFX ← sticky

if (guard | sticky) then
    SPEFSCRFINXS ← 1

// Round the integer result
if ((round = ROUND) & (SPEFSCRFINXE = 0)) then
    if (SPEFSCRFRMC = 0b00) then // nearest
        if (guard) then
            if (sticky | (result & 0x00000001)) then
                result ← result + 1
        else if ((SPEFSCRFRMC & 0b10) = 0b10) then // infinity modes
            // implementation dependent

```



```

if (signed = SIGN) then
  if (fpsign = 1) then
    result ← -result + 1
return result

```

### 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
  SignalFPErr(LOWER, FINV)
  if (Isa64NaN(fp)) then
    return 0x00000000 // all NaNs
  if (signed = SIGN) then
    if (fpsign = 1) then
      return 0x80000000
    else
      return 0x7fffffff
  else
    if (fpsign = 1) then
      return 0x00000000
    else
      return 0xffffffff

if (Isa64Denorm(fp)) then
  SignalFPErr(LOWER, FINV)
  return 0x00000000 // regardless of sign

if ((signed = UNSIGN) & (fpsign = 1)) then
  SignalFPErr(LOWER, FOVF) // overflow
  return 0x00000000

if ((fpexp = 0) & (fpfrac = 0)) then
  return 0x00000000 // all zero values

if (fractional = I) then // convert to integer
  max_exp ← 1054
  shift ← 1054 - fpexp
  if (signed ← SIGN) then
    if ((fpexp ≠ 1054) | (fpfrac ≠ 0) | (fpsign ≠ 1)) then
      max_exp ← max_exp - 1
else // fractional conversion
  max_exp ← 1022
  shift ← 1022 - fpexp
  if (signed = SIGN) then
    shift ← shift + 1

```

```

if (fpexp > max_exp) then
  SignalFPErr(LOWER, FOVF) // overflow
  if (signed = SIGN) then
    if (fpsign = 1) then
      return 0x80000000
    else
      return 0x7fffffff
  else
    return 0xffffffff

result ← 0b1 || fpfrac[0:30] // add U to frac
guard ← fpfrac[31]
sticky ← (fpfrac[32:63] ≠ 0)
for (n ← 0; n < shift; n ← n + 1) do
  sticky ← sticky | guard
  guard ← result & 0x00000001
  result ← result > 1

// Report sticky and guard bits

SPEFSCRFG ← guard
SPEFSCRFX ← sticky

if (guard | sticky) then
  SPEFSCRFINXS ← 1

// Round the result

if ((round = ROUND) & (SPEFSCRFINXE = 0)) then
  if (SPEFSCRFRMC = 0b00) then // nearest
    if (guard) then
      if (sticky | (result & 0x00000001)) then
        result ← result + 1
    else if ((SPEFSCRFRMC & 0b10) = 0b10) then // infinity modes
      // implementation dependent

if (signed = SIGN) then
  if (fpsign = 1) then
    result ← -result + 1

return result

```

### 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
  SignalFPErr(LOWER, FINV)
  if (Isa64NaN(fp)) then
    return 0x00000000_00000000 // all NaNs

```

```

if (signed = SIGN) then
  if (fpsign = 1) then
    return 0x80000000_00000000
  else
    return 0x7fffffff_ffffff
else
  if (fpsign = 1) then
    return 0x00000000_00000000
  else
    return 0xffffffff_ffffff

if (Isa64Denorm(fp)) then
  SignalFPErr(LOWER, FINV)
  return 0x00000000_00000000 // regardless of sign

if ((signed = UNSIGN) & (fpsign = 1)) then
  SignalFPErr(LOWER, FOVF) // overflow
  return 0x00000000_00000000

if ((fpexp = 0) & (fpfrac = 0)) then
  return 0x00000000_00000000 // all zero values

max_exp ← 1086
shift ← 1086 - fpexp
if (signed = SIGN) then
  if ((fpexp ≠ 1086) | (fpfrac ≠ 0) | (fpsign ≠ 1)) then
    max_exp ← max_exp - 1

if (fpexp > max_exp) then
  SignalFPErr(LOWER, FOVF) // overflow
  if (signed = SIGN) then
    if (fpsign = 1) then
      return 0x80000000_00000000
    else
      return 0x7fffffff_ffffff
  else
    return 0xffffffff_ffffff

result ← 0b1 || fpfrac || 0b000000000000 // add U to frac
guard ← 0
sticky ← 0
for (n ← 0; n < shift; n ← n + 1) do
  sticky ← sticky | guard
  guard ← result & 0x00000000_00000001
  result ← result > 1

// Report sticky and guard bits

SPEFSCRFG ← guard
SPEFSCRFX ← sticky

if (guard | sticky) then
  SPEFSCRFINXS ← 1

// Round the result

```

```

if ((round = ROUND) & (SPEFSCRFINXE = 0)) then
  if (SPEFSCRFRMC = 0b00) then // nearest
    if (guard) then
      if (sticky | (result & 0x00000000_00000001)) then
        result ← result + 1
      else if ((SPEFSCRFRMC & 0b10) = 0b10) then // infinity modes
        // implementation dependent
    if (signed = SIGN) then
      if (fpsign = 1) then
        result ← -result + 1
  return result

```

### Convert to single-precision floating-point from integer word with saturation

```

// Convert from integer/fractional 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)
FP32format result;
resultsign ← 0
if (v = 0) then
  result ← 0
  if (upper_lower = UPPER) then
    SPEFSCRFGH ← 0
    SPEFSCRFXH ← 0
  else
    SPEFSCRFG ← 0
    SPEFSCRFX ← 0
else
  if (signed = SIGN) then
    if (v0 = 1) then
      v ← -v + 1
      resultsign ← 1
  if (fractional = F) then // fractional bit pos alignment
    maxexp ← 127
    if (signed = UNSIGN) then
      maxexp ← maxexp - 1
  else
    maxexp ← 158 // integer bit pos alignment
  sc ← 0
  while (v0 = 0)
    v ← v << 1
    sc ← sc + 1
  v0 ← 0 // clear U bit
  resultexp ← maxexp - sc
  guard ← v24
  sticky ← (v25:31 ≠ 0)

```

```

// Report sticky and guard bits
if (upper_lower = UPPER) then
    SPEFSCRFGH ← guard
    SPEFSCRFXH ← sticky
else
    SPEFSCRFG ← guard
    SPEFSCRFX ← sticky

if (guard | sticky) then
    SPEFSCRFINXS ← 1

// Round the result

resultfrac ← v1:23
result ← Round32(result, guard, sticky)

return result

```

### Convert to double-precision floating-point from integer word with saturation

```

// Convert from integer/fractional to 64-bit floating point
// signed = SIGN or UNSIGN
// fractional = F (fractional) or I (integer)
CnvtI32ToFP64Sat(v, signed, fractional)
FP64format result;
resultsign ← 0
if (v = 0) then
    result ← 0
    SPEFSCRFG ← 0
    SPEFSCRFX ← 0
else
    if (signed = SIGN) then
        if (v[0] = 1) then
            v ← -v + 1
            resultsign ← 1
    if (fractional = F) then // fractional bit pos alignment
        maxexp ← 1023
        if (signed = UNSIGN) then
            maxexp ← maxexp - 1
    else
        maxexp ← 1054 // integer bit pos alignment
    sc ← 0
    while (v0 = 0)
        v ← v << 1
        sc ← sc + 1
    v0 ← 0 // clear U bit
    resultexp ← maxexp - sc

// Report sticky and guard bits
SPEFSCRFG ← 0
SPEFSCRFX ← 0

```

```

    resultfrac ← v1:31 || 210
    return result

```

### 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;
resultsign ← 0
if (v = 0) then
    result ← 0
    SPEFSCRFG ← 0
    SPEFSCRFX ← 0
else
    if (signed = SIGN) then
        if (v0 = 1) then
            v ← ¬v + 1
            resultsign ← 1
        maxexp ← 1054
        sc ← 0
        while (v0 = 0)
            v ← v << 1
            sc ← sc + 1
        v0 ← 0 // clear U bit
        resultexp ← maxexp - sc
        guard ← v53
        sticky ← (v54:63 ≠ 0)
// Report sticky and guard bits
    SPEFSCRFG ← guard
    SPEFSCRFX ← sticky
    if (guard | sticky) then
        SPEFSCRFINXS ← 1
// Round the result
    resultfrac ← v1:52
    result ← Round64(result, guard, sticky)
    return result

```

### 7.3.3 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
    
```

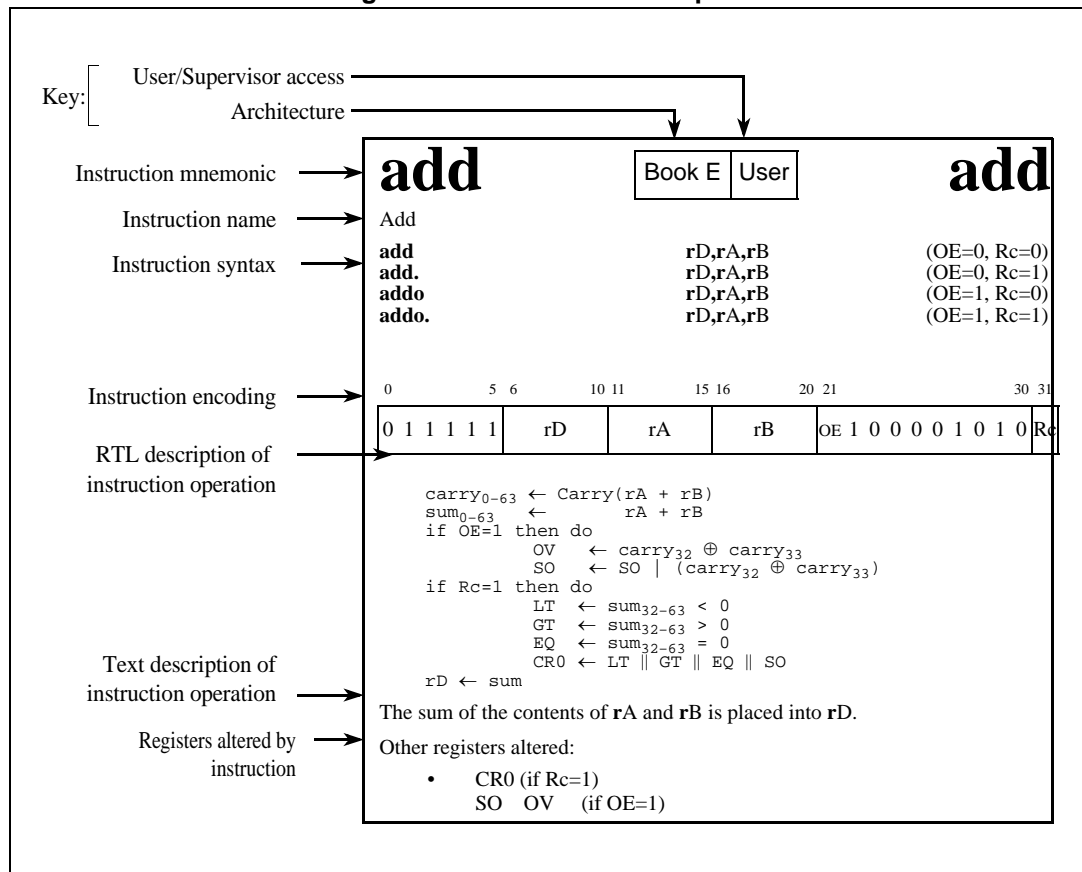
### 7.3.4 Embedded floating-point results

[Appendix E: Embedded floating-point results](#) , summarizes results of various types of SPE and SPFP floating-point operations on various combinations of input operands.

## 7.4 Instruction set


The rest of this chapter describes individual instructions, which are listed in alphabetical order by mnemonic. [Figure 86](#) shows the format for instruction description pages.

**Figure 86. Instruction description**



*Note:* The execution unit that executes the instruction may not be the same for all processors.

<b>add</b>	Book E	User	<b>add</b>
Add			
<b>add</b>	<b>rD,rA,rB</b>		(OE=0, Rc=0)
<b>add.</b>	<b>rD,rA,rB</b>		(OE=0, Rc=1)
<b>addo</b>	<b>rD,rA,rB</b>		(OE=1, Rc=0)
<b>addo.</b>	<b>rD,rA,rB</b>		(OE=1, Rc=1)



```

carry0:63 ← Carry(rA + rB)
sum0:63 ← rA + rB
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← sum
carry0:63 ← Carry(rA + rB)
sum0:63 ← rA + rB
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← sum
    
```


The sum of the contents of rA and rB is placed into rD.

Other registers altered:

- CR0 (if Rc=1)
- SO OV (if OE=1)

<b>_addx</b>	VLE	User	<b>addx</b>
Add			
<b>se_add</b>	<b>rX, rY</b>		






$$\text{sum}_{32:63} \leftarrow \text{GPR}(\text{RX}) + \text{GPR}(\text{RY})$$

$$\text{GPR}(\text{RX}) \leftarrow \text{sum}_{32:63}$$

The sum of the contents of GPR(rX) and the contents of GPR(rY) is placed into GPR(rX).

Special Registers Altered: None

<b>addc</b>	Book E   User	<b>addc</b>
<b>Add Carrying</b>		
<b>add</b>	rD,rA,rB	(OE=0, Rc=0)
<b>add.</b>	rD,rA,rB	(OE=0, Rc=1)
<b>addo</b>	rD,rA,rB	(OE=1, Rc=0)
<b>addo.</b>	rD,rA,rB	(OE=1, Rc=1)



$$\text{carry}_{0:63} \leftarrow \text{Carry}(\text{rA} + \text{rB})$$

$$\text{sum}_{0:63} \leftarrow \text{rA} + \text{rB}$$

if OE=1 then do

$$\text{OV} \leftarrow \text{carry}_{32} \oplus \text{carry}_{33}$$

$$\text{SO} \leftarrow \text{SO} | (\text{carry}_{32} \oplus \text{carry}_{33})$$

if Rc=1 then do

$$\text{LT} \leftarrow \text{sum}_{32:63} < 0$$

$$\text{GT} \leftarrow \text{sum}_{32:63} > 0$$

$$\text{EQ} \leftarrow \text{sum}_{32:63} = 0$$

$$\text{CR0} \leftarrow \text{LT} || \text{GT} || \text{EQ} || \text{SO}$$

$$\text{rD} \leftarrow \text{sum}$$


$$\text{CA} \leftarrow \text{carry}_{32}$$

The sum of the contents of rA and rB is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)

<b>adde</b>	Book E	User	<b>adde</b>
<b>Add Extended</b>			
<b>adde</b>	<b>rD,rA,rB</b>		(OE=0, Rc=0)
<b>adde.</b>	<b>rD,rA,rB</b>		(OE=0, Rc=1)
<b>addeo</b>	<b>rD,rA,rB</b>		(OE=1, Rc=0)
<b>addeo.</b>	<b>rD,rA,rB</b>		(OE=1, Rc=1)



```


if E=0 then Cin ← CA
carry0:63 ← Carry(rA + rB + Cin)
sum0:63 ← rA + rB + Cin
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← sum
CA ← carry32
    
```

For **adde[o][.]**, the sum of the contents of rA, the contents of rB, and CA is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)

**addi** Book E User **addi**  
**Add immediate [shifted]**  
**addi** rD,rA,SIMM (S=0)  
**addis** rD,rA,SIMM (S=1)



if rA=0 then a ← 640 else a ← rA  
 if s=0 then b ← EXTS(SIMM)  
 if s=1 then b ← EXTS(SIMM || 160)  
 rD ← a + b

If **addi** and rA=0, the sign-extended value of the SIMM field is placed into rD.

If **addi** and rA≠0, the sum of the contents of rA and the sign-extended value of field SIMM is placed into rD.


If **addis** and rA=0, the sign-extended value of the SIMM field, concatenated with 16 zeros, is placed into rD.

If **addis** and rA≠0, the sum of the contents of rA and the sign-extended value of the SIMM field concatenated with 16 zeros, is placed into rD.

Other registers altered: None

**\_addix** VLE User **\_addix**  
**Add [2 operand] Immediate [Shifted][and Record]**

**e\_add16i** rD,rA,SI




a ← GPR(RA)  
 b ← EXTS(SI)  
 GPR(RD) ← a + b

The sum of the contents of GPR(rA) and the sign-extended value of field SI is placed into GPR(rD).

Special Registers Altered: None

**e\_add2i** rA,SI




```

SI ← SI0:4 || SI5:15
sum32:63 ← GPR(RA) + EXTS(SI)
LT ← sum32:63 < 0
GT ← sum32:63 > 0
EQ ← sum32:63 = 0
CR0 ← LT || GT || EQ || SO
GPR(RA) ← sum32:63
    
```

The sum of the contents of GPR(rA) and the sign-extended value of SI is placed into GPR(rA).

Special Registers Altered: CR0

**e\_add2is** rA,SI




```

SI ← SI0:4 || SI5:15
sum32:63 ← GPR(RD) + (SI || 16'0)
GPR(RA) ← sum32:63
    
```

The sum of the contents of GPR(rA) and the value of SI concatenated with 16 zeros is placed into GPR(rAarav2006).

Special Registers Altered: None

**e\_addi** rD,rA,SCI8 (Rc=0)  
**e\_addi.** rD,rA,SCI8 (Rc=1)




```

imm ← SCI8(F,SCL,UI8)
sum32:63 ← GPR(RA) + imm
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RD) ← sum32:63
    
```

The sum of the contents of GPR(rA) and the value of SCI8 is placed into GPR(rD).

Special Registers Altered: CR0 (if Rc = 1)

**se\_addi** rX,OIMM



1. OIMM = OIM5 + 1.




$$GPR(RX) \leftarrow GPR(RX) + (2^7 0 \parallel OFFSET(OIM5))$$

The sum of the contents of GPR(rX) and the zero-extended offset value of OIM5 (a final value in the range 1–32), is placed into GPR(rX).

Special Registers Altered: None

<b>addic</b>	Book E	User	<b>addic</b>
<b>Add immediate carrying [and record]</b>			
<b>addic</b>	rD,rA,SIMM		(Rc=0)
<b>addic.</b>	rD,rA,SIMM		(Rc=1)



```


carry0:63 ← Carry(rA + EXTS(SIMM))
sum0:63 ← rA + EXTS(SIMM)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← rA+EXTS(SIMM)
CA ← carry32
    
```

The sum of the contents of rA and the sign-extended value of the SIMM field is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)

<b>_addicx</b>	VLE	User	<b>_addicx</b>
<b>Add immediate Carrying [and record]</b>			
<b>e_addic</b>	rD,rA,SCI8		(Rc=0)
<b>e_addic.</b>	rD,rA,SCI8		(Rc=1)



```

imm ← SCI8(F,SCL,UI8)
carry32:63 ← Carry(GPR(RA) + imm)
sum32:63 ← GPR(RA) + imm
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
    
```

GPR(RD) ← sum<sub>32:63</sub>  
 CA ← carry<sub>32</sub>

The sum of the contents of GPR(rA) and the value of SCI8 is placed into GPR(rD).


Special Registers Altered: CA, CR0 (if Rc=1)

**addme**

Book E	User
--------	------

**addme**  
**Add to minus one extended**

<b>addme</b>	rD,rA	(OE=0, Rc=0)
<b>addme.</b>	rD,rA	(OE=0, Rc=1)
<b>addmeo</b>	rD,rA	(OE=1, Rc=0)
<b>addmeo.</b>	rD,rA	(OE=1, Rc=1)



```

if E=0 then Cin ← CA
carry0:63 ← Carry(rA + Cin + 0xFFFF_FFFF_FFFF_FFFF)
sum0:63 ← rA + Cin + 0xFFFF_FFFF_FFFF_FFFF
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)

if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO

rD ← sum
CA ← carry32
    
```

For **addme[o][.]**, the sum of the contents of rA, CA, and <sup>64</sup>1 is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)


**addze**

Book E	User
--------	------

**addze**

**Add to zero extended**

<b>addze</b>	rD,rA	(OE=0, Rc=0)
<b>addze.</b>	rD,rA	(OE=0, Rc=1)
<b>addzeo</b>	rD,rA	(OE=1, Rc=0)
<b>addzeo.</b>	rD,rA	(OE=1, Rc=1)



```

if E=0 then Cin ← CA
carry0:63 ← Carry(rA + Cin)
sum0:63 ← rA + Cin
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)

if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO

rD ← sum
CA ← carry32
    
```

For **addze[o][.]**, the sum of the contents of rA and CA is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)


**and**

Book E	User
--------	------


**and**

**AND [Immediate [Shifted] | with Complement]**


<b>and</b>	rA,rS,rB	(Rc=0)
<b>and.</b>	rA,rS,rB	(Rc=1)



**andi**  $rA, rS, UIMM$  (S=0, Rc=1)  
**andis.**  $rA, rS, UIMM$  (S=1, Rc=1)



**andc**  $rA, rS, rB$  (Rc=0)  
**andc.**  $rA, rS, rB$  (Rc=1)



```

if 'andi.' then b ← 480 || UIMM
if 'andis.' then b ← 320 || UIMM || 160
if 'and[.]' then b ← rB
if 'andc[.]' then b ← ¬rB
result0:63 ← rS & b
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
rA ← result
    
```


For **andi.**, the contents of rS are ANDed with <sup>48</sup>0 || UIMM.  
 For **andis.**, the contents of rS are ANDed with <sup>32</sup>0 || UIMM || <sup>16</sup>0.  
 For **and[.]**, the contents of rS are ANDed with the contents of rB.  
 For **andc[.]**, the contents of rS are ANDed with the one's complement of the contents of rB.  
 The result is placed into rA.  
 Other registers altered: CR0 (if Rc=1)

**\_andx**

VLE	User
-----	------


**\_andx**  
**AND [2 operand] [Immediate | with Complement] [and Record]**

**se\_and**  $rX, rY$  (Rc=0)  
**se\_and.**  $rX, rY$  (Rc=1)




**e\_and2i**  $rD, U$






**e\_and2is** **rD,U**




**e\_andi** **rA,rS,SCI8** **(Rc=0)**


**e\_andi.** **rA,rS,SCI8** **(Rc=1)**



**e\_andi** **rX,UI5**



**se\_andc** **rX,rY**



```

if 'e_andi[.]' then b ← SCI8(F,SCL,UI8)
if 'se_andi' then b ← UI5
if 'se_and[.]' then b ← GPR(RY)
if 'se_andc' then b ← ¬GPR(RY)
if 'e_and2i.' then b ← 160 || UI0:4 || UI5:15
if 'e_and2is.' then b ← UI0:4 || UI5:15 || 160
result32:63 ← GPR(RS or RD or RX) & b
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
if 'se_and[cij]' then GPR(RX) ← result32:63 else GPR(RA or RD) ← result32:63
    
```

For **e\_andi[.]**, the contents of GPR(rS) are ANDed with the value of SCI8.

For **e\_and2i.**, the contents of GPR(rD) are ANDed with <sup>16</sup>0 || UI.

For **e\_and2is.**, the contents of GPR(rD) are ANDed with UI || <sup>16</sup>0.

For **se\_andi**, the contents of GPR(rX) are ANDed with the value of UI5.

For **se\_and[.]**, the contents of GPR(rX) are ANDed with the contents of GPR(rY).

For **se\_andc**, the contents of GPR(rX) are ANDed with the one's complement of the contents of GPR(rY).

The result is placed into GPR(rA) or GPR(rX) (**se\_and[ic][.]**)


Special Registers Altered: CR0 (if Rc = 1)

**b**

Book E	User
--------	------

**b**  
**Branch [and Link] [Absolute]**

<b>b</b>	LI	(AA=0, LK=0)
<b>ba</b>	LI	(AA=1, LK=0)
<b>bl</b>	LI	(AA=0, LK=1)
<b>bla</b>	LI	(AA=1, LK=1)



if AA=1 then a ← <sup>64</sup>0 else a ← CIA  
 if E=0 then NIA ← <sup>32</sup>0 || (a + EXTS(LI||0b00))<sub>32:63</sub>  
 if LK=1 then LR ← CIA + 4

The branch target effective address (BTEA) is calculated as follows:

- For 32-bit implementations, BTEA is bits 32–63 of the sum of the current instruction address (CIA), or 32 zeros if AA=1, and the sign-extended value of the LI instruction field concatenated with 0b00

BTEA is the address of the next instruction to be executed.

If LK=1, the sum CIA+4 is placed into the LR.


Other registers altered: LR (if LK=1)

**\_bx**

VLE	User
-----	------

**\_bx**  
**Branch [and Link]**

<b>e_b</b>	BD24	(LK=0)
<b>e_bx</b>	BD24	(LK=1)



a ← CIA  
 NIA ← (a + EXTS(BD24||0b0))<sub>32:63</sub>  
 if LK=1 then LR ← CIA + 4

Let the BTEA be calculated as follows:

- For **e\_b[!]**, let BTEA be the sum of the CIA and the sign-extended value of the BD24 instruction field concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA+4 is placed into the LR.

Special Registers Altered: LR (if LK = 1)

**se\_b** BD8 (LK=0)  
**se\_bl** BD8 (LK=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	1	1	0	1	0	0	LK	BD8							

$a \leftarrow CIA$   
 $NIA \leftarrow (a + EXTS(BD8||0b0))_{32:63}$   
 if LK=1 then  $LR \leftarrow CIA + 2$

Let the BTEA be calculated as follows:

- For **se\_b**[i], let BTEA be the sum of the CIA and the sign-extended value of the BD8 instruction field concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA+2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)

**bc**

Book E	User
--------	------

**bc**  
**Branch conditional [and link] [absolute]**

**bc** BO,BI,BD (AA=0, LK=0)  
**bca** BO,BI,BD (AA=1, LK=0)  
**bcl** BO,BI,BD (AA=0, LK=1)  
**bcla** BO,BI,BD (AA=1, LK=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	0	0	BO			BI			BD												AA	LK						

if  $\neg BO_2$  then  $CTR_{32:63} \leftarrow CTR_{32:63} : 1$   
 $ctr\_ok \leftarrow BO_2 | ((CTR_{32:63} \neq 0) \oplus BO_3)$   
 $cond\_ok \leftarrow BO_0 | (CR_{BI+32} \equiv BO_1)$   
 if  $ctr\_ok \ \& \ cond\_ok$  then  
     if AA=1 then  $a \leftarrow {}^{64}0$  else  $a \leftarrow CIA$   
     if E=0 then  $NIA \leftarrow {}^{32}0 || (a + EXTS(BD||0b00))_{32:63}$   
 else  $NIA \leftarrow CIA + 4$   
 if LK=1 then  $LR \leftarrow CIA + 4$

The branch target effective address (BTEA) is calculated as follows:

- For 32-bit implementations, BTEA is bits 32–63 of the sum of the current instruction address (CIA), or 32 zeros if AA=1, and the sign-extended value of the LI instruction field concatenated with 0b00

The BO instruction field specifies any conditions that must be met for the branch to be taken, as defined in [Section 4.3.1.12: Conditional branch control](#). The sum BI+32 specifies the CR bit to be used.

The BI field specifies the CR bit used as the condition of the branch, as shown in [Table 200](#).

**Table 200. BI operand settings for CR fields**


CRn Bits	CR Bits	BI	Description
CR0[0]	32	00000	Negative (LT)—Set when the result is negative.
CR0[1]	33	00001	Positive (GT)—Set when the result is positive (and not zero).
CR0[2]	34	00010	Zero (EQ)—Set when the result is zero.
CR0[3]	35	00011	Summary overflow (SO). Copy of XER[SO] at the instruction's completion.
CR1[0]	36	00100	Copy of FPSCR[FX] at the instruction's completion.
CR1[1]	37	00101	Copy of FPSCR[FEX] at the instruction's completion.
CR1[2]	38	00110	Copy of FPSCR[VX] at the instruction's completion.
CR1[3]	39	00111	Copy of FPSCR[OX] at the instruction's completion.
CRn[0]	40	01000	Less than or floating-point less than (LT, FL). For integer compare instructions: rA < SIMM or rB (signed comparison) or rA < UIMM or rB (unsigned comparison). For floating-point compare instructions: frA < frB.
	44	01100	
	48	10000	
	52	10100	
	56	11000	
CRn[1]	60	11100	
	41	01001	Greater than or floating-point greater than (GT, FG). For integer compare instructions: rA > SIMM or rB (signed comparison) or rA > UIMM or rB (unsigned comparison). For floating-point compare instructions: frA > frB.
	45	01101	
	49	10001	
	53	10101	
57	11001		
CRn[2]	61	11101	
	42	01010	Equal or floating-point equal (EQ, FE). For integer compare instructions: rA = SIMM, UIMM, or rB. For floating-point compare instructions: frA = frB.
	46	01110	
	50	10010	
	54	10110	
58	11010		
CRn[3]	62	11110	
	43	01011	Summary overflow or floating-point unordered (SO, FU). For integer compare instructions, this is a copy of XER[SO] at the completion of the instruction. For floating-point compare instructions, one or both of frA and frB is a NaN.
	47	01111	
	51	10011	
	55	10111	
59	11011		
	63	11111	

If the branch conditions are met, the BTEA is the address of the next instruction to be executed.

If LK=1, the sum CIA + 4 is placed into the LR.


Other registers altered:

- CTR(if BO<sub>2</sub>=0)  
LR(if LK=1)



**Branch Conditional [and Link]**

**e\_bc** BO32,BI32,BD15 (LK=0)  
**e\_bcl** BO32,BI32,BD15 (LK=1)



```

if BO320 then CTR32:63 ← CTR32:63 - 1
ctr_ok ← ¬BO320 | ((CTR32:63 ≠ 0) ⊕ BO321)
cond_ok ← BO320 | (CRBI32+32 ≡ BO321)
if ctr_ok & cond_ok then
    NIA ← (CIA + EXTS(BD15 || 0b0))32:63
else
    NIA ← CIA + 4
if LK=1 then LR ← CIA + 4
    
```

Let the BTEA be calculated as follows:

- For **e\_bc**[I], let BTEA be the sum of the CIA and the sign-extended value of the BD15 instruction field concatenated with 0b0.


BO32 specifies any conditions that must be met for the branch to be taken, as defined in [Section 13.2.2: Branch instructions](#). The sum BI32+32 specifies the CR bit. Only CR[32–47] may be specified.

If the branch conditions are met, the BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA + 4 is placed into the LR.

Special Registers Altered: CTR (if BO32<sub>0</sub> = 1)  
LR (if LK = 1)

**se\_bc** BO16,BI16,BD8



```

cond_ok ← (CRBI16+32 ≡ BO16)
if cond_ok then
    NIA ← (CIA + EXTS(BD8 || 0b0))32:63
else
    NIA ← CIA + 2
    
```

Let the BTEA be calculated as follows:

- For **se\_bc**, BTEA is the sum of the CIA and the sign-extended value of the BD8 instruction field concatenated with 0b0.

BO16 specifies any conditions that must be met for the branch to be taken, as defined in [Section 13.2.2: Branch instructions](#). The sum BI16+32 specifies CR bit; only CR[32–35] may be specified.

If the branch conditions are met, the BTEA is the address of the next instruction to be executed.

Special Registers Altered: None

<b>bcctr</b>	Book E	User	<b>bcctr</b>
<b>Branch conditional to count register [and link]</b>			
<b>bcctr</b>	BO,BI		(LK=0)
<b>bcctrl</b>	BO,BI		(LK=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	0	0	1	1	BO		BI		///			1		0		0		0		1		0		0		0		0		LK	

```

cond_ok ← BO0 | (CRBI+32 ≡ BO1)
if cond_ok & E=0 then NIA ← 320 || CTR32:61 || 0b00
if ¬cond_ok then NIA ← CIA + 4
if LK=1 then LR ← CIA + 4
    
```

The branch target effective address (BTEA) is calculated as follows:

- For **bcctr[I]**, BTEA is the contents of CTR[32–61] concatenated with 0b00.

BO specifies conditions that must be met for the branch to be taken. BI+32 specifies the CR bit to be used; see [Table 201](#).

**Table 201. BI operand settings for CR fields**

CRn Bits	CR Bits	BI	Description
CR0[0]	32	00000	Negative (LT)—Set when the result is negative.
CR0[1]	33	00001	Positive (GT)—Set when the result is positive (and not zero).
CR0[2]	34	00010	Zero (EQ)—Set when the result is zero.
CR0[3]	35	00011	Summary overflow (SO). Copy of XER[SO] at the instruction's completion.
CR1[0]	36	00100	Copy of FPSCR[FX] at the instruction's completion.
CR1[1]	37	00101	Copy of FPSCR[FEX] at the instruction's completion.
CR1[2]	38	00110	Copy of FPSCR[VX] at the instruction's completion.
CR1[3]	39	00111	Copy of FPSCR[OX] at the instruction's completion.

**Table 201. BI operand settings for CR fields (continued)**

CRn Bits	CR Bits	BI	Description
CRn[0]	40	01000	Less than or floating-point less than (LT, FL). For integer compare instructions: rA < SIMM or rB (signed comparison) or rA < UIMM or rB (unsigned comparison). For floating-point compare instructions: frA < frB.
	44	01100	
	48	10000	
	52	10100	
	56	11000	
	60	11100	
CRn[1]	41	01001	Greater than or floating-point greater than (GT, FG). For integer compare instructions: rA > SIMM or rB (signed comparison) or rA > UIMM or rB (unsigned comparison). For floating-point compare instructions: frA > frB.
	45	01101	
	49	10001	
	53	10101	
	57	11001	
	61	11101	
CRn[2]	42	01010	Equal or floating-point equal (EQ, FE). For integer compare instructions: rA = SIMM, UIMM, or rB. For floating-point compare instructions: frA = frB.
	46	01110	
	50	10010	
	54	10110	
	58	11010	
	62	11110	
CRn[3]	43	01011	Summary overflow or floating-point unordered (SO, FU). For integer compare instructions, this is a copy of XER[SO] at the completion of the instruction. For floating-point compare instructions, one or both of frA and frB is a NaN.
	47	01111	
	51	10011	
	55	10111	
	59	11011	
	63	11111	

If the condition is met, the BTEA is the address of the next instruction to be executed.

If LK=1, the sum CIA + 4 is placed into the LR.

If the decrement and test CTR option is specified (BO[2]=0), the instruction form is invalid.

Other registers altered: LR (if LK=1)

<b>bclr</b>	Book E	User	<b>bclr</b>
<b>Branch conditional to link register [and link]</b>			
<b>bclr</b>	BO,BI		(LK=0)
<b>bclrl</b>	BO,BI		(LK=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	1	0	0	1	1	BO		BI		///			0		0		0		0		0		1		0		0		0		0		LK

$$\text{if } \neg \text{BO}_2 \text{ then } \text{CTR}_{32:63} \leftarrow \text{CTR}_{32:63} - 1$$

$$\text{ctr\_ok} \leftarrow \text{BO}_2 \mid ((\text{CTR}_{32:63} \neq 0) \oplus \text{BO}_3)$$

```

cond_ok ← BO0 | (CRBI+32 ≡ BO1)
if ctr_ok & cond_ok & E=0 then NIA ← 320 || LR32:61 || 0b00
if ¬(ctr_ok & cond_ok) then NIA ← CIA + 4
if LK=1 then LR ← CIA + 4
    
```

The branch target effective address (BTEA) is calculated as follows:

- For **bclr**[I], BTEA is the contents of LR[32–61] concatenated with 0b00.

The BO field specifies any conditions that must be met for the branch to be taken, as defined in [Section 4.3.1.12: Conditional branch control](#). The sum BI+32 specifies the CR bit to be used.

The BI field specifies the CR bit used as the condition of the branch, as shown in [Table 202](#).

**Table 202. BI operand settings for CR fields**

CRn Bits	CR Bits	BI	Description
CR0[0]	32	00000	Negative (LT)—Set when the result is negative.
CR0[1]	33	00001	Positive (GT)—Set when the result is positive (and not zero).
CR0[2]	34	00010	Zero (EQ)—Set when the result is zero.
CR0[3]	35	00011	Summary overflow (SO). Copy of XER[SO] at the instruction's completion.
CR1[0]	36	00100	Copy of FPSCR[FX] at the instruction's completion.
CR1[1]	37	00101	Copy of FPSCR[FEX] at the instruction's completion.
CR1[2]	38	00110	Copy of FPSCR[VX] at the instruction's completion.
CR1[3]	39	00111	Copy of FPSCR[OX] at the instruction's completion.
CRn[0]	40	01000	Less than or floating-point less than (LT, FL). For integer compare instructions: <b>rA &lt; SIMM</b> or <b>rB</b> (signed comparison) or <b>rA &lt; UIMM</b> or <b>rB</b> (unsigned comparison). For floating-point compare instructions: <b>frA &lt; frB</b> .
	44	01100	
	48	10000	
	52	10100	
	56	11000	
CRn[1]	60	11100	
	41	01001	Greater than or floating-point greater than (GT, FG). For integer compare instructions: <b>rA &gt; SIMM</b> or <b>rB</b> (signed comparison) or <b>rA &gt; UIMM</b> or <b>rB</b> (unsigned comparison). For floating-point compare instructions: <b>frA &gt; frB</b> .
	45	01101	
	49	10001	
	53	10101	
57	11001		
CRn[2]	61	11101	
	42	01010	Equal or floating-point equal (EQ, FE). For integer compare instructions: <b>rA = SIMM</b> , <b>UIMM</b> , or <b>rB</b> . For floating-point compare instructions: <b>frA = frB</b> .
	46	01110	
	50	10010	
	54	10110	
58	11010		
	62	11110	



**Table 202. BI operand settings for CR fields (continued)**

CRn Bits	CR Bits	BI	Description
CRn[3]	43	01011	Summary overflow or floating-point unordered (SO, FU).
	47	01111	
	51	10011	For integer compare instructions, this is a copy of XER[SO] at the completion of the instruction.
	55	10111	For floating-point compare instructions, one or both of frA and frB is a NaN.
	59	11011	
	63	11111	

If the condition is met, the BTEA is the address of the next instruction to be executed.


If LK=1, the sum CIA + 4 is placed into the LR.

Other registers altered:

- CTR (if BO<sub>2</sub>=0)
- LR (if LK=1)

<b>_bclri</b>	<table border="0" style="width: 100%;"> <tr> <td style="width: 50%; text-align: center; padding: 2px;">VLE</td> <td style="width: 50%; text-align: center; padding: 2px;">User</td> </tr> </table>	VLE	User	<b>_bclri</b>
VLE	User			

**Bit Clear Immediate**



$$\begin{aligned}
 a &\leftarrow \text{UI5} \\
 b &\leftarrow {}^a1 \parallel 0 \parallel {}^{31-a}1 \\
 \text{result}_{32:63} &\leftarrow \text{GPR(RX)} \& b \\
 \text{GPR(RX)} &\leftarrow \text{result}_{32:63}
 \end{aligned}$$


For **se\_bclri**, the bit of GPR(rX) specified by the value of UI5 is cleared and all other bits in GPR(rX) remain unaffected.

Special Registers Altered: None

<b>_bctrx</b>	<table border="0" style="width: 100%;"> <tr> <td style="width: 50%; text-align: center; padding: 2px;">VLE</td> <td style="width: 50%; text-align: center; padding: 2px;">User</td> </tr> </table>	VLE	User	<b>_bctrx</b>
VLE	User			

**Branch to Count Register [and Link]**

**se\_bctr** (LK = 0)  
**se\_bctrl** (LK = 1)



$$\begin{aligned}
 \text{NIA} &\leftarrow \text{CTR}_{32:62} \parallel 0b0 \\
 \text{if LK=1 then LR} &\leftarrow \text{CIA} + 2
 \end{aligned}$$

Let the BTEA be calculated as follows:

- For **se\_bctrl**[I], let BTEA be bits 32–62 of the contents of the CTR concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.



If LK = 1, the sum CIA + 2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)

**\_bgeni** **\_bgeni**

VLE	User
-----	------

**Bit Generate Immediate**

**se\_bgeni** rX,UI5

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	1	1	0	0	0	1	UI5					RX				

$a \leftarrow \text{UI5}$   
 $b \leftarrow {}^a0 \parallel 1 \parallel {}^{31-a}0$   
 $\text{GPR(RX)} \leftarrow b$

For **se\_bgeni**, a constant value consisting of a single '1' bit surrounded by '0's is generated and the value is placed into GPR(RX). The position of the '1' bit is specified by the UI5 field.

Special Registers Altered: None

**\_blrx** **\_blrx**

VLE	User
-----	------

**Branch to Link Register [and Link]**

**se\_blr** (LK = 0)  
**se\_brlr** (LK = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
															LK

$\text{NIA} \leftarrow \text{LR}_{32:62} \parallel 0b0$   
 if LK=1 then  $\text{LR} \leftarrow \text{CIA} + 2$

Let the BTEA be calculated as follows:

- For **se\_blr[I]**, let BTEA be bits 32–62 of the contents of the LR concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA + 2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)

**\_bmaski** **\_bmaski**

VLE	User
-----	------

**Bit Mask Generate Immediate**

**se\_bmaski** rX,UI5

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	1	0	1	1	0	UI5					RX				

$a \leftarrow \text{UI5}$   
 if  $a = 0$  then  $b \leftarrow {}^{32}1$  else  $b \leftarrow {}^{32-a}0 \parallel a1$   
 $\text{GPR(RX)} \leftarrow b$


For **se\_bmaski**, a constant value consisting of a mask of low-order '1' bits that is zero-extended to 32 bits is generated, and the value is placed into GPR(**rX**). The number of low-order '1' bits is specified by the UI5 field. If UI5 is 0b00000, a value of all '1's is generated

Special Registers Altered: None

**brinc** SPE APU User **brinc**

**Bit reversed increment**

**brinc** **rD,rA,rB**



```

n ← MASKBITS // Imp dependent # of mask bits
mask ← rB64-n:63 // Least sig. n bits of register
a ← rA64-n:63
d ← bitreverse(1 + bitreverse(a | (~ mask)))
rD ← rA0:63-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. **rB** 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
    
```

**rB** 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  $\log_2 n$  1s in the least significant bit positions and 0s 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 1s are shifted left by  $\log_2$  (size of the data) and 0s are placed in the least significant bit positions. [Table 203](#) shows example values of masks for different data sizes and number of data.

**Table 203. Data samples and sizes**


Number of data samples	Data size			
	Byte	Half word	Word	Double word
8	000...00000111	000...00001110	000...000011100	000...0000111000
16	000...00001111	000...00011110	000...000111100	000...0001111000
32	000...00011111	000...00111110	000...001111100	000...0011111000
64	000...00111111	000...01111110	000...011111100	000...0111111000

**\_bseti** VLE User **\_bseti**

**Bit Set Immediate**

**se\_bseti** **rX,UI5**





$a \leftarrow \text{UI5}$   
 $b \leftarrow a_0 \parallel 1 \parallel a_{31-a_0}$   
 $\text{result}_{32:63} \leftarrow \text{GPR}(\text{RX}) \mid b$   
 $\text{GPR}(\text{RX}) \leftarrow \text{result}_{32:63}$


For **se\_bseti**, the bit of GPR(rX) specified by the value of UI5 is set, and all other bits in GPR(rX) remain unaffected.

Special Registers Altered: None



**Bit Test Immediate**


**se\_btsti** **rX,UI5**



$a \leftarrow \text{UI5}$   
 $b \leftarrow a_0 \parallel 1 \parallel a_{31-a_0}$   
 $c \leftarrow \text{GPR}(\text{RX}) \& b$   
if  $c = 0$  then  $d \leftarrow 0b001$  else  $d \leftarrow 0b010$   
 $\text{CR}_{0:3} \leftarrow d \parallel \text{XER}_{\text{SO}}$


For **se\_btsti**, the bit of GPR(rX) specified by the value of UI5 is tested for equality to '1'. The result of the test is recorded in the CR. EQ is set if the tested bit is clear, LT is cleared, and GT is set to the inverse value of EQ.

Special Registers Altered: CR[0–3]




**Compare [immediate]**

**cmp** **crD,L,rA,rB**



**cmpi** **crD,L,rA,SIMM**



if  $L=0$  then  $a \leftarrow \text{EXTS}(\text{rA}_{32:63})$   
else  $a \leftarrow \text{rA}$   
if 'cmpi' then  $b \leftarrow \text{EXTS}(\text{SIMM})$   
if 'cmp' &  $L=0$  then  $b \leftarrow \text{EXTS}(\text{rB}_{32:63})$   
if 'cmp' &  $L=1$  then  $b \leftarrow \text{rB}$



if a < b then c ← 0b100  
 if a > b then c ← 0b010  
 if a = b then c ← 0b001  
 $CR_{4 \times crD+32:4 \times crD+35} \leftarrow c \parallel XER_{SO}$

If **cmp** and L=0, the contents of rA[32–63] are compared with the contents of rB[32–63], treating the operands as signed integers.

If **cmpi** and L=0, the contents of rA[32–63] are compared with the sign-extended value of the SIMM field, treating the operands as signed integers.


The result of the comparison is placed into CR field **crD**.

Other registers altered: CR field **crD**


**\_cmp** VLE    User **cmp**

**Compare [Immediate]**

**e\_cmp16i**                      rA,SI



**e\_cmpi**                                      crD32,rA,SCI8



a ← GPR(RA)<sub>32:63</sub>  
 if 'e\_cmpi' then b ← SCI8(F,SCL,UI8)  
 if 'e\_cmp16i' then b ← EXT(SI<sub>0:4</sub> || SI<sub>5:15</sub>)  
 if a < b then c ← 0b100  
 if a > b then c ← 0b010  
 if a = b then c ← 0b001  
 if 'e\_cmpi' then  $CR_{4 \times crD32+32:4 \times crD32+35} \leftarrow c \parallel XER_{SO}$  // only CR0-CR3  
 if 'e\_cmp16i' then  $CR_{32:35} \leftarrow c \parallel XER_{SO}$  // only CR0


If **e\_cmpi**, GPR(rA) contents are compared with the value of SCI8, treating operands as signed integers.

If **e\_cmp16i**, GPR(rA) contents are compared with the sign-extended value of the SI field, treating operands as signed integers.


The result of the comparison is placed into CR field **crD** (**crD32**). For **e\_cmpi**, only CR0–CR3 may be specified. For **e\_cmp16i**, only CR0 may be specified.

Special Registers Altered: CR field **crD** (**crD32**) (CR0 for **e\_cmp16i**)

**se\_cmp**                                      rX,rY



**se\_cmpi**                                      rX,UI5



```

a ← GPR(RX)32:63
if 'se_cmpi' then b ← 270 || UI5
if 'se_cmp' then b ← GPR(RY)32:63
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR0:3 ← c || XERSO
    
```

If **se\_cmp**, the contents of GPR(**rX**) are compared with the contents of GPR(**rY**), treating the operands as signed integers. The result of the comparison is placed into CR field 0.

If **se\_cmpi**, the contents of GPR(**rX**) are compared with the value of the zero-extended UI5 field, treating the operands as signed integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]


**\_cmph**

VLE	User
-----	------

**\_cmph**

**Compare Halfword [Immediate]**

**e\_cmph** **crD,rA,rB**




```

a ← EXTS(GPR(RA)48:63)
b ← EXTS(GPR(RB)48:63)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR4×CRD+32:4×CRD+35 ← c || XERSO
    
```

For **e\_cmph**, the contents of the low-order 16 bits of GPR(**rA**) and GPR(**rB**) are compared, treating the operands as signed integers. The result of the comparison is placed into CR field CRD.

Special Registers Altered: CR field CRD

**se\_cmph** **rX,rY**



```

a ← EXTS(GPR(RX)48:63)
b ← EXTS(GPR(RY)48:63)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR0:3 ← c || XERSO
    
```

For **se\_cmph**, the contents of the low-order 16 bits of GPR(**rX**) and GPR(**rY**) are compared, treating the operands as signed integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]

**e\_cmph16i** **rA,SI**



0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	0	SI <sub>0:4</sub>				RA				1	0	1	1	0	SI <sub>5:15</sub>												

```
a ← EXTS(GPR(RA)48:63)
b ← EXTS(SI0:4 || SI5:15)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR32:35 ← c || XERSO // only CR0
```

The contents of the lower 16-bits of GPR(rA) are sign-extended and compared with the sign-extended value of the SI field, treating the operands as signed integers.

The result of the comparison is placed into CR0.

Special Registers Altered: CR0

<b>_cmphi</b>	VLE	User	<b>_cmphi</b>
<b>_cmphi _cmphi _cmphi _cmphi</b>			

### Compare Halfword Logical [Immediate]

**e\_cmphi**                      crD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRD	/	RA				RB				0	0	0	0	1	0	1	1	1	0	/					

```
a ← EXTZ(GPR(RA)48:63)
b ← EXTZ(GPR(RB)48:63)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR4×CRD+32:4×CRD+35 ← c || XERSO
```

For **e\_cmphi**, the contents of the low-order 16 bits of GPR(rA) and GPR(rB) are compared, treating the operands as unsigned integers. The result of the comparison is placed into CR field CRD.

Special Registers Altered: CR field CRD

**se\_cmphi**                      rX,rY

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	1	1	1	1	RY				RX			


```
a ← GPR(RX)48:63
b ← GPR(RY)48:63
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR0:3 ← c || XERSO
```

For **se\_cmphi**, the contents of the low-order 16 bits of GPR(rX) and GPR(rY) are compared, treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0-3]



**e\_cmphl16i** **rA,UI**



```

a ← 160 || GPR(RA)48:63
b ← 160 || UI0:4 || UI5:15
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR32:35 ← c || XERSO // only CR0
    
```

The contents of the lower 16-bits of GPR(**rA**) are zero-extended and compared with the zero-extended value of the UI field, treating the operands as unsigned integers.


The result of the comparison is placed into CR0.

Special Registers Altered: CR0


**cmpl** Book E | User **cmpl**

**Compare logical [immediate]**

**cmpl** **crD,L,rA,rB**



**cmpli** **crD,L,rA,UIMM**



```

if L=0 then a ← 320 || rA32:63
else      a ← rA
if 'cmpli' then b ← 480 || UIMM
if 'cmpl' & L=0 then b ← 320 || rB32:63
if 'cmpl' & L=1 then b ← rB
if a <u b then c ← 0b100
if a >u b then c ← 0b010
if a = b then c ← 0b001
CR4×crD+32:4×crD+35 ← c || XERSO
    
```

If **cmpl** and L=0, the contents of **rA**[32–63] are compared with the contents of **rB**[32–63], treating the operands as unsigned integers.

If **cmpli** and L=0, the contents of **rA**[32–63] are compared with the zero-extended value of the UIMM field, treating the operands as unsigned integers.

The result of the comparison is placed into CR field **crD**.


Other registers altered: CR field **crD**




**\_cimpl** \_cimpl  
 \_cimpl \_cimpl \_cimpl \_cimpl VLE User

**Compare Logical [Immediate]**

**e\_cimpl16i** rA,UI



**e\_cmpli** crD32,rA,SCI8



$a \leftarrow \text{GPR}(RA)_{32:63}$   
 if 'e\_cmpli' then  $b \leftarrow \text{SCI8}(F,SCL,UI8)$   
 if 'e\_cmpl16i' then  $b \leftarrow 160 \parallel UI_{0:4} \parallel UI_{5:15}$   
 if  $a <_u b$  then  $c \leftarrow 0b100$   
 if  $a >_u b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 if 'e\_cmpli' then  $CR_{4 \times CRD32+32:4 \times CRD32+35} \leftarrow c \parallel XER_{SO}$  // only CR0-CR3  
 if 'e\_cmpl16i' then  $CR_{32:35} \leftarrow c \parallel XER_{SO}$  // only CR0

If **e\_cmpli**, the contents of bits 32–63 of GPR(rA) are compared with the value of SCI8, treating the operands as unsigned integers.


L must be 0 for 32-bit implementations

If **e\_cmpl16i**, the contents of GPR(rA) are compared with the zero-extended value of the UI field, treating the operands as unsigned integers.


The result of the comparison is placed into CR field CRD (CRD32). For **e\_cmpli**, only CR0–CR3 may be specified. For **e\_cmpl16i**, only CR0 may be specified.

Special Registers Altered: CR field CRD (CRD32) (CR0 for **e\_cmpl16i**)

**se\_cmpl** rX,rY



**se\_cmpli** rX,OIMM



1. OIMM = OIM5 + 1

$a \leftarrow \text{GPR}(RX)_{32:63}$   
 if 'se\_cmpli' then  $b \leftarrow 270 \parallel \text{OFFSET}(OIM5)$   
 if 'se\_cmpl' then  $b \leftarrow \text{GPR}(RY)_{32:63}$   
 if  $a <_u b$  then  $c \leftarrow 0b100$   
 if  $a >_u b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $CR_{0:3} \leftarrow c \parallel XER_{SO}$

If **se\_cmpl**, the contents of GPR(**rX**) are compared with the contents of GPR(**rY**), treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.

If **se\_cmpli**, the contents of GPR(**rX**) are compared with the value of the zero-extended offset value of the OIM5 field (a final value in the range 1–32), treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.


Special Registers Altered: CR[0–3]

**cntlzw** Book E User **cntlzw**

**Count leading zeros (word)**

**cntlzw** **rA,rS**(Z=0, Rc=0)

**cntlzw.** **rS**(Z=0, Rc=1)



```

if 'cntlzd' then n ← 0 else n ← 32
i ← 0
do while n < 64
  if rSn = 1 then leave
  n ← n + 1
  i ← i + 1
rA ← i
if Rc=1 then do
  GT ← i > 0
  EQ ← i = 0
  CR0 ← 0b0 || GT || EQ || SO
    
```


For **cntlzw[.]**, a count of the number of consecutive zero bits starting at **rS**[32] is placed into **rA**. This number ranges from 0 to 32, inclusive. If **Rc=1**, CR field 0 is set to reflect the result.

Other registers altered: CR0 (if **Rc=1**)

**crand** Book E User **crand**

**Condition register AND**

**crand** **crbD,crbA,crbB**



$$CR_{crbD+32} \leftarrow CR_{crbA+32} \& CR_{crbB+32}$$

The content of bit **crbA+32** of CR is ANDed with the content of bit **crbB+32** of CR, and the result is placed into bit **crbD+32** of CR.

Other registers altered: CR

**\_crand** **\_crand**  
**\_crand \_crand \_crand \_crand** VLE User

**Condition Register AND**

**e\_crاند** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	CRBD					CRBA					CRBB					0	1	0	0	0	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \& CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ANDed with the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

Condition Register AND with Complement

**e\_crاندc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	CRBD					CRBA					CRBB					0	0	1	0	0	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \& \neg CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ANDed with the one's complement of the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

CR Equivalent

**e\_creqv** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	CRBD					CRBA					CRBB					0	1	0	0	1	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \equiv CR_{BB+32}$$

The content of bit CRBA+32 of the CR is XORed with the content of bit CRBB+32 of the CR, and the one's complement of result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**crاندc** **crاندc**  
Book E User

**Condition register AND with complement**

**crاندc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	0	0	1	1	crbD					crbA					crbB					0	0	1	0	0	0	0	0	0	0	1	/

$$CR_{crbD+32} \leftarrow CR_{crbA+32} \& \neg CR_{crbB+32}$$

The content of bit **crbA+32** of CR is ANDed with the one's complement of the content of bit **crbB+32** of CR, and the result is placed into bit **crbD+32** of CR.

Other registers altered: CR

**creqv** Book E User **creqv**  
**Condition register equivalent**

**creqv** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>			<b>crbA</b>			<b>crbB</b>			0	1	0	0	1	0	0	0	0	0	1	/					

$$CR_{crbD+32} \leftarrow CR_{crbA+32} \oplus CR_{crbB+32}$$

The content of bit **crbA + 32** of CR is XORed with the content of bit **crbB + 32** of CR, and the one's complement of result is placed into bit **crbD+32** of CR.

Other registers altered: CR

**crnand** Book E User **crnand**  
**Condition register NAND**

**crnand** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>			<b>crbA</b>			<b>crbB</b>			0	0	1	1	1	0	0	0	0	1	/						

$$CR_{crbD+32} \leftarrow \neg(CR_{crbA+32} \& CR_{crbB+32})$$

The content of bit **crbA+32** of CR is ANDed with the content of bit **crbB+32** of CR, and the one's complement of the result is placed into bit **crbD+32** of CR.

Other registers altered: CR

**\_crnand** VLE User **\_crnand**  
**\_crnand \_crnand \_crnand \_crnand**  
**Condition Register NAND**

**e\_crnand** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRBD			CRBA			CRBB			0	0	1	1	1	0	0	0	0	1	/						

$$CR_{BT+32} \leftarrow \neg(CR_{BA+32} \& CR_{BB+32})$$

The content of bit CRBA+32 of the CR is ANDed with the content of bit CRBB+32 of the CR, and the one's complement of the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**crnor** Book E User **crnor**

**Condition register NOR**

**crnor** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>			<b>crbA</b>			<b>crbB</b>			0	0	0	0	1	0	0	0	0	0	1	/					

$$CR_{crbD+32} \leftarrow \neg(CR_{crbA+32} \mid CR_{crbB+32})$$

The content of bit **crbA**+32 of CR is ORed with the content of bit **crbB**+32 of CR, and the one's complement of the result is placed into bit **crbD**+32 of CR.

Other registers altered: CR

**\_crnor** VLE User **\_crnor**

**Condition Register NOR**

**e\_crnor** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRBD			CRBA			CRBB			0	0	0	0	1	0	0	0	0	1	/						

$$CR_{BT+32} \leftarrow \neg(CR_{BA+32} \mid CR_{BB+32})$$

The content of bit CRBA+32 of the CR is ORed with the content of bit CRBB+32 of the CR, and the one's complement of the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**cror** Book E User **cror**

**Condition register OR**

**cror** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>			<b>crbA</b>			<b>crbB</b>			0	1	1	1	0	0	0	0	0	1	/						

$$CR_{crbD+32} \leftarrow CR_{crbA+32} \mid CR_{crbB+32}$$

The content of bit **crbA**+32 of CR is ORed with the content of bit **crbB**+32 of CR, and the result is placed into bit **crbD**+32 of CR.

Other registers altered: CR

**\_cror** VLE User **\_cror**

**Condition Register OR**



**e\_cror** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	CRBD					CRBA					CRBB					0	1	1	1	0	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \mid CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ORed with the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**crorc** Book E User **crorc**

**Condition register OR with complement**

**crorc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>					<b>crbA</b>					<b>crbB</b>					0	1	1	0	1	0	0	0	0	1	/

$$CR_{crbD+32} \leftarrow CR_{crbA+32} \mid \neg CR_{crbB+32}$$

The content of bit **crbA**+32 of CR is ORed with the one's complement of the content of bit **crbB**+32 of CR, and the result is placed into bit **crbD**+32 of CR.

Other registers altered: CR

**\_crorc** VLE User **\_crorc**

**\_crorc \_crorc \_crorc \_crorc**

**Condition Register OR with Complement**

**e\_crorc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRBD					CRBA					CRBB					0	1	1	0	1	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \mid \neg CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ORed with the one's complement of the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**crxor** Book E User **crxor**

**Condition register XOR**


**crxor** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1	<b>crbD</b>					<b>crbA</b>					<b>crbB</b>					0	0	1	1	0	0	0	0	0	1	/

$$CR_{\text{crbD}+32} \leftarrow CR_{\text{crbA}+32} \oplus CR_{\text{crbB}+32}$$


The content of bit **crbA+32** of CR is XORed with the content of bit **crbB+32** of CR, and the result is placed into bit **crbD+32** of CR.

Other registers altered: CR



**Condition Register XOR**


**e\_crxor**                                      **crbD,crbA,crbB**



$$CR_{\text{crbD}+32} \leftarrow CR_{\text{crbA}+32} \oplus CR_{\text{crbB}+32}$$


The content of bit **crbA+32** of the CR is XORed with the content of bit **crbB+32** of the CR, and the result is placed into bit **crbD+32** of the CR.

Special Registers Altered: CR



**Data cache block allocate**

**dcba**    **rA,rB**



$$\text{if } rA=0 \text{ then } a \leftarrow 64_0 \text{ else } a \leftarrow rA$$
$$EA \leftarrow 32_0 \parallel (a + rB)_{32:63}$$
$$\text{AllocateDataCacheBlock}(EA)$$

EA calculation:                      Addressing ModeEA for rA=0EA for rA≠0

$$32_0 \parallel rB_{32:63} \quad 32_0 \parallel (rA+rB)_{32:63}$$

**dcba** is a hint that performance would likely improve if the block containing the byte addressed by EA is established in the data cache without fetching the block from main memory, because the program is likely to soon store into a portion of the block and the contents of the rest of the block are not meaningful to the program. If the hint is honored, the contents of the block are undefined when the instruction completes. The hint is ignored if the block is caching-inhibited.

If the block containing the byte addressed by EA is in memory that is memory-coherence required and the block exists in a data cache of any other processors, it is kept coherent in those caches.

This instruction is treated as a storeexcept that an interrupt is not taken for a translation or protection violation.



This instruction may establish a block in the data cache without verifying that the associated real address is valid. This can cause a delayed machine check interrupt.

Other registers altered: None

**dcbf** Book E User **dcbf**  
**Data cache block flush**

**dcbf** rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	///						rA	rB	0	0	0	1	0	1	0	1	1	0	/							

if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)<sub>32:63</sub>  
 FlushDataCacheBlock( EA )

EA calculation:      Addressing ModeEA for rA=0EA for rA≠0  
                             320 || rB<sub>32:63</sub> 320 || (rA+rB)<sub>32:63</sub>

If the block containing the byte addressed by EA is in memory that is memory-coherence required, a block containing the byte addressed by EA is in the data cache of any processor, and any locations in the block are considered to be modified there, then those locations are written to main memory. Additional locations in the block may also be written to main memory. The block is invalidated in the data caches of all processors.

If the block containing the byte addressed by EA is in memory that is not memory-coherence required, a block containing the byte addressed by EA is in the data cache of this processor and any locations in the block are considered to be modified there, then those locations are written to main memory. Additional locations in the block may also be written to main memory. The block is invalidated in the data cache of this processor.

On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

The function of this instruction is independent of whether the block containing the byte addressed by EA is in memory that is write-through required or caching-inhibited.

This instruction is treated as a load. See [Section 4.6.6.8: Cache management instructions](#).

Other registers altered: None

**dcbi** Book E User **dcbi**  
**Data cache block invalidate**

**dcbi** rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	///						rA	rB	0	1	1	1	0	1	0	1	1	0	/							



if  $rA=0$  then  $a \leftarrow 640$  else  $a \leftarrow rA$   
 $EA \leftarrow 320 \parallel (a + rB)_{32:63}$   
 InvalidateDataCacheBlock( EA )  
 EA calculation: Addressing ModeEA for  $rA=0$ EA for  $rA \neq 0$   
 $320 \parallel rB_{32:63} \quad 320 \parallel (rA+rB)_{32:63}$

If the block containing the byte addressed by EA is in coherence-required memory and any block containing the addressed byte is any processors' data cache is invalidated in those caches. On some implementations, before the block is invalidated, if any locations in the block are considered to be modified in any such data cache, those locations are written to main memory and additional locations in the block may be written to main memory.

If the block containing the byte addressed by EA is not coherence-required memory and a block containing the byte addressed by EA is in the data cache of this processor, then the block is invalidated in that data cache. On some implementations, before the block is invalidated, any locations in the block considered modified in that data cache are written to main memory; additional locations in the block may be written to main memory.

**dcbi** is treated as a store on implementations that invalidate a block without first writing to main memory all locations in the block that are considered to be modified in the data cache, except that the invalidation is not ordered by **mbar**. On other implementations this instruction is treated as a load.

Additional information about this instruction is as follows.

- The data cache block size for **dcbi** is the same as for **dcbf**.
- If a processor holds a reservation and some other processor executes a **dcbi** to the same reservation granule, whether the reservation is lost is undefined.

Other registers altered: None

**dcblc**

Cache locking APU	User
-------------------	------

**dcblc**

**Data cache block lock clear**

**dcblc** CT,rA,rB Form: X

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1					CT				rA					rB			0	1	1	0	0	0	0	1	1	0	/

if  $rA = 0$  then  $a \leftarrow 640$  else  $a \leftarrow GPR(rA)$   
 if Mode32 then  $EA \leftarrow 320 \parallel (a + GPR(rB))_{32:63}$   
 if Mode64 then  $EA \leftarrow a + GPR(rB)$   
 DataCacheBlockClearLock(CT, EA)  
 EA calculation: EA for  $rA=0$ EA for  $rA \neq 0$   
 $320 \parallel GPR(rB)_{32:63} \quad 320 \parallel$   
 $(GPR(rA) + GPR(rB))_{32:63}$

The data cache specified by CT has the cache line corresponding to EA unlocked allowing the line to participate in the normal replacement policy.

Cache lock clear instructions remove locks previously set by cache lock set instructions.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

The instruction is treated as a load with respect to translation and memory protection and can cause DSI and DTLB error interrupts accordingly.

An unable-to-unlock condition is said to occur any of the following conditions exist:

- The target address is marked cache-inhibited, or the storage attributes of the address uses a coherency protocol that does not support locking.
- The target cache is disabled or not present.
- The CT field of the instructions contains a value not supported by the implementation.
- The target address is not in the cache or is present in the cache but is not locked.

If an unable-to-unlock condition occurs, no cache operation is performed.

**EIS Specifics**

Clearing and then setting L1CSR0[CLFR] allows system software to clear all L1 data cache locking bits without knowing the addresses of the lines locked.

**dcbst** Book E User **dcbst**

**Data Cache Block Store**

**dcbst** rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	///						rA				rB				0	0	0	0	1	1	0	1	1	0	/	

if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)32:63  
 StoreDataCacheBlock( EA )  
 EA calculation: Addressing ModeEA for rA=0EA for rA≠0  
 $320 \mid \mid rB_{32:63} \mid \mid 320 \mid \mid (rA+rB)_{32:63}$

If the block containing the byte addressed by EA is in memory that is memory-coherence required and a block containing the byte addressed by EA is in the data cache of any processor, and any locations in the block are considered to be modified there, those locations are written to main memory. Additional locations in the block may be written to main memory. The block ceases to be considered to be modified in that data cache.

If the block containing the byte addressed by EA is in memory that is not memory-coherence required and a block containing the byte addressed by EA is in the data cache of this processor and any locations in the block are considered to be modified there, those locations are written to main memory. Additional locations in the block may be written to main memory. The block ceases to be considered to be modified in that cache.

The function of this instruction is independent of whether the block containing the byte addressed by EA is in memory that is write-through required or caching-inhibited.

This instruction is treated as a load.

On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

Other registers altered: None

**dcbt** Book E User **dcbt**

**Data cache block touch**



**dcbt** CT,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CT					rA					rB					0	1	0	0	0	1	0	1	1	0	/

if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)32:63  
 PrefetchDataCacheBlock( CT, EA )  
 EA calculation: Addressing ModeEA for rA=0EA for rA≠0  
 320 || rB32:63 320 || (rA+rB) 32:63

*Section 4.3.1.18: User-level cache instructions*, lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

Implementations should perform no operation when CT specifies a value not supported by the implementation.

The hint is ignored if the block is caching-inhibited.

This instruction is treated as a load except that an interrupt is not taken for a translation or protection violation.

Other registers altered: None

<b>dcbtls</b>	Cache locking APU	User	<b>dcbtls</b>
---------------	-------------------	------	---------------

**Data cache block touch and lock set**

**dcbtls** CT,rA,rB Form: X

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CT					rA					rB					0	0	1	0	1	0	0	1	1	0	/

if rA = 0 then a ← 640 else a ← GPR(rA)  
 if Mode32 then EA ← 320 || (a + GPR(rB))32:63  
 if Mode64 then EA ← a + GPR(rB)  
 PrefetchDataCacheBlockLockSet(CT, EA)  
 EA calculation: EA for rA=0EA for rA≠0  
 320 || GPR(rB) 32:63 320 ||  
 (GPR(rA) + GPR(rB)) 32:63

The data cache specified by CT has the cache line corresponding to EA loaded and locked into the cache. If the line already exists in the cache, it is locked without being refetched.

Cache touch and lock set instructions let software lock cache lines into the cache to provide lower latency for critical cache accesses and more deterministic behavior. Locked lines do not participate in the normal replacement policy when a line must be victimized for replacement.

*Section 4.3.1.18: User-level cache instructions*, lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

The instruction is treated as a load with respect to translation and memory protection and can cause DSI and DTLB error interrupts accordingly.



An unable to lock condition is said to occur any of the following conditions exist:

- The target address is marked cache-inhibited, or the storage attributes of the address uses a coherency protocol that does not support locking.
- The target cache is disabled or not present.
- The CT field of the instructions contains a value not supported by the implementation.

If an unable to lock condition occurs, no cache operation is performed and LICSR0[DCUL] is set appropriately.

Overlocking is said to exist is all available ways for a given cache index are already locked. If overlocking occurs for **dcbtIs** and if the lock was targeted for the primary cache (CT = 0), the requested line is not locked into the cache. When overlock occurs, L1CSR1[DCLO] is set. If L1CSR1[DCLOA] is set, the requested line is locked into the cache and implementation dependent line currently locked in the cache is evicted.

The results of overlocking and unable to lock conditions for caches other than the primary cache and secondary cache are defined as part of the architecture for the specific cache hierarchy designated by CT.

Other registers altered:

- L1CSR0[DCUL] if unable to lock occurs
- L1CSR0[DCLO] (L2CSR[L2CLO]) if lock overflow occurs


**dcbtst**

Book E	User
--------	------

**dcbtst**

**Data cache block touch for store**

**dcbtst** CT,rA,rB



if rA=0 then a ← <sup>64</sup>0 else a ← rA  
 EA ← <sup>32</sup>0 || (a + rB)<sub>32:63</sub>  
 PrefetchForstoreDataCacheBlock( CT, EA )  
 EA calculation:            Addressing ModeEA for rA=0EA for rA≠0  
                                   <sup>32</sup>0 | | rB<sub>32:63</sub> <sup>32</sup>0 || (rA+rB)<sub>32:63</sub>

If CT=0, this instruction is a hint that performance would likely be improved if the block containing the byte addressed by EA is fetched into the data cache, because the program will probably soon store into the addressed byte.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

Implementations should perform no operation when CT specifies a value not supported by the implementation.

The hint is ignored if the block is caching-inhibited.

This instruction is treated as a load , except that an interrupt is not taken for a translation or protection violation.

Other registers altered: None



**dcbtstls**

Cache locking APU	User
-------------------	------

**dcbtstls**

**Data cache block touch for store and lock set**

**dcbtstls**

CT,rA,rB

Form: X

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	1	1	1	1	1																															/

if rA = 0 then a ← <sup>64</sup>0 else a ← GPR(rA)  
 if Mode32 then EA ← <sup>32</sup>0 || (a + GPR(rB))<sub>32:63</sub>  
 if Mode64 then EA ← a + GPR(rB)  
 PrefetchDataCacheBlockLockSet(CT, EA)  
 EA calculation:           EA for rA=0EA for rA≠0  
                                   <sup>32</sup>0 || GPR(rB)<sub>32:63</sub> <sup>32</sup>0 || (GPR(rA) + GPR(rB))<sub>32:63</sub>

The data cache specified by CT has the cache line corresponding to EA loaded and locked into the cache. If the line already exists in the cache, it is locked without refetching from memory.

Cache touch and lock set instructions allow software to lock lines into the cache to shorten latency for critical cache accesses and more deterministic behavior. Lines locked in the cache do not participate in the normal replacement policy when a line must be victimized for replacement.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

[Table 114](#) describes how this instruction is treated with respect to translation and memory protection.

For unable-to-lock conditions, described in [Section 9.1.1.4: Unable-to-lock conditions](#), no cache operation is performed and L1CSR0[DCUL] is set.

Overlocking occurs when all available ways for a given cache index are already locked. If an overlocking condition occurs for a **dcbtstls** instruction and if the lock was targeted for the primary cache or secondary cache (CT = 0 or CT = 2), the requested line is not locked into the cache. When overlock occurs, L1CSR1[DCLO] (L2CSR[L2CLO] for CT = 2) is set. If L1CSR1[DCLOA] is set (or L2CSR[L2CLOA] for CT = 2), the requested line is locked into the cache and implementation dependent line currently locked in the cache is evicted. If system software wants to precisely determine if an overlock event has occurred in the L1 data cache, it must perform the following code sequence:

```
dcbtstls
msync
mfspr (L1CSR0)
(check L1CSR0[DCUL] bit for data cache index unable-to-lock condition)
(check L1CSR0[DCLO] bit for data cache index overlock condition)
```

Results of overlocking and unable-to-lock conditions for caches other than the primary and secondary cache are defined as part of the architecture for the cache hierarchy designated by CT.

Other registers altered:

- L1CSR0[DCUL] if unable to lock occurs
- L1CSR0[DCLO] (L2CSR[L2CLO]) if lock overflow occurs

**EIS specifics:**

Clearing and then setting L1CSR0[CLFR] allows system software to clear all data cache locking bits without knowing the addresses of the lines locked.


**dcbz**

Book E	User
--------	------

**dcbz**

**Data cache block set to zero**

**dcbzr**The data cache specified by CT has theA,rB



```

if rA=0 then a ← 640 else a ← rA
EA ← 320 || (a + rB)32:63
ZeroDataCacheBlock( EA )
EA calculation:      Addressing ModeEA for rA=0EA for rA≠0
                    320 || rB32:63 320 || (rA+rB) 32:63
    
```

If the block containing the addressed byte is in the data cache, all bytes of the block are cleared.

If the block containing the byte addressed by EA is not in the data cache and is in memory that is not caching-inhibited, the block is established in the data cache without fetching the block from main memory, and all bytes of the block are cleared.

If the block containing the byte addressed by EA is not in the data cache and is in storage that is not caching inhibited and cannot be established in the cache, then one of the following occurs:

- All bytes of the area of main storage that corresponds to the addressed block are set to zero
- An alignment interrupt is taken

If the block containing the byte addressed by EA is in storage that is caching inhibited or write through required, one of the following occurs:

- All bytes of the area of main storage that corresponds to the addressed block are set to zero
- An alignment interrupt is taken.

If the block containing the byte addressed by EA is in memory-coherence required memory and the block exists in any other processors' data cache, it is kept coherent in those caches.

**dcbz** may establish a block in the data cache without verifying that the associated real address is valid. This can cause a delayed machine check interrupt.

**dcbz** is treated as a store.

- On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.
- **dcbz** may cause a cache-locking exception on some implementations. See the user documentation.

Other registers altered: None




Programming note: If the block containing the byte addressed by EA is in memory that is caching-inhibited or write-through required, the alignment interrupt handler should clear all bytes of the area of main memory that corresponds to the addressed block.

**divw** Book E User **divw**

**Divide word**

<b>divw</b>	rD,rA,rB	(OE=0, Rc=0)
<b>divw.</b>	rD,rA,rB	(OE=0, Rc=1)
<b>divwo</b>	rD,rA,rB	(OE=1, Rc=0)
<b>divwo.</b>	rD,rA,rB	(OE=1, Rc=1)



```

dividend0:31 ← rA32:63
divisor0:31 ← rB32:63
quotient0:31 ← dividend ÷ divisor
if OE=1 then do
    OV ← ( (rA32:63 = -231) & (rB32:63 = -1) ) | (rB32:63 = 0)
    SO ← SO | OV
if Rc=1 then do
    LT ← quotient < 0
    GT ← quotient > 0
    EQ ← quotient = 0
    CR0 ← LT || GT || EQ || SO

rD32:63 ← quotient
rD0:31 ← undefined
    
```

The 32-bit quotient of the contents of rA[32–63] divided by the contents of rB[32–63] is placed into rD[32–63]. rD[0–31] are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as signed integers. The quotient is the unique signed integer that satisfies the following:

$$\text{dividend} = (\text{quotient} \times \text{divisor}) + r$$

Here,  $0 \leq r < |\text{divisor}|$  if the dividend is nonnegative and  $-|\text{divisor}| < r \leq 0$  if it is negative.

If any of the following divisions is attempted, the contents of rD are undefined as are (if Rc=1) the contents of the CR0[LT,GT,EQ]. In these cases, if OE=1, OV is set.

0x8000\_0000 ÷ -1  
 <anything> ÷ 0

Other registers altered:

- CR0 (if Rc=1)
- SO OV (if OE=1)


**divwu** Book E User **divwu**

**Divide word unsigned**

<b>divwu</b>	rD,rA,rB	(OE=0, Rc=0)
<b>divwu.</b>	rD,rA,rB	(OE=0, Rc=1)



**divwuo**                                 rD,rA,rB                                 (OE=1, Rc=0)  
**divwuo.**                                 rD,rA,rB                                 (OE=1, Rc=1)



```

dividend0:31 ← rA32:63
divisor0:31 ← rB32:63
quotient0:31 ← dividend ÷ divisor
if OE=1 then do
    OV ← (rB32:63=0)
    SO ← SO | OV
if Rc=1 then do
    LT ← quotient < 0
    GT ← quotient > 0
    EQ ← quotient = 0
    CR0 ← LT || GT || EQ || SO

rD32:63 ← quotient
rD0:31 ← undefined
    
```

The 32-bit quotient of the contents of rA[32–63] divided by the contents of rB[32–63] is placed into rD[32–63]. rD[0–31] are undefined. The remainder is not supplied as a result.

Both operands and the quotient are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR field 0 are set by signed comparison of the result to zero. The quotient is the unique unsigned integer that satisfies the following:

$$\text{dividend} = (\text{quotient} \times \text{divisor}) + r$$

Here,  $0 \leq r < \text{divisor}$ .

If an attempt is made to perform the following division, the contents of rD are undefined as are (if Rc=1) the contents of the LT, GT, and EQ bits of CR0. In this case, if OE=1 OV is set.

$$\langle \text{anything} \rangle \div 0$$


Other registers altered:

- CR0 (if Rc=1)
- SO OV (if OE=1)

**efdabs**                                 Scalar DFPF APU | User                                 **efdabs**

**Floating-point double-precision absolute value**

**efdabs**                                 rD,rA



$$rD_{0:63} \leftarrow 0b0 \parallel rA_{1:63}$$

The sign bit of rA is set to 0 and the result is placed into rD.

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 **rA** 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** Scalar DFP APU | User **efdadd**

**Floating-point double-precision add**

**efdadd** **rD,rA,rB**



$$rD_{0:63} \leftarrow rA_{0:63} +_{dp} rB_{0:63}$$

**rA** is added to **rB** and the result is stored in **rD**. If **rA** is NaN or infinity, the result is either *pmax* ( $a_{sign}=0$ ), or *nmax* ( $a_{sign}=1$ ). Otherwise, if **rB** is NaN or infinity, the result is either *pmax* ( $b_{sign}=0$ ), or *nmax* ( $b_{sign}=1$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in **rD**. 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 **rA** or **rB** 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** Scalar DFP APU | User **efdcfs**

**Floating-point double-precision convert from single-precision**

**efdcfs** **rD,rB**



```

FP32format f;
FP64format result;
f ← rB32:63
if (fexp = 0) & (ffrac = 0) then
    result ← fsign || 630 // signed zero value
else if Isa32NaNorInfinity(f) | Isa32Denorm(f) then
    SPEFSCRFINV ← 1
    result ← fsign || 0b1111111110 || 521 // max value
else if Isa32Denorm(f) then
    SPEFSCRFINV ← 1
    result ← fsign || 630
else
    resultsign ← fsign
    resultexp ← fexp - 127 + 1023
    resultfrac ← ffrac || 290

rD0:63 = result
    
```

The single-precision floating-point value in the low element of **rB** is converted to a double-precision floating-point value and the result is placed into **rD**. The rounding mode is not used since this conversion is always exact.

Exceptions:

If the low element of **rB** 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.

*Note: Architecture Note: This instruction is optional if neither the embedded scalar single-precision floating-point APU or the embedded vector single-precision floating-point APU are implemented.*


**efdcfsf**

Scalar DPFAPU	User
---------------	------

**efdcfsf**

**Convert floating-point double-precision from signed fraction**

**efdcfsf** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtI32ToFP64}(rB_{32:63}, \text{SIGN}, F)$$

The signed fractional low element in **rB** is converted to a double-precision floating-point value using the current rounding mode and the result is placed into **rD**.

Exceptions:

None.


**efdcfsi**

Scalar DPFAPU	User
---------------	------

**efdcfsi**

**Convert floating-point double-precision from signed integer**

**efdcfsi** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtSI32ToFP64}(rB_{32:63}, \text{SIGN}, I)$$

The signed integer low element in **rB** is converted to a double-precision floating-point value using the current rounding mode and the result is placed into **rD**.

Exceptions:

None.


**efdcfsid** 

Scalar DPFP APU	User
-----------------	------

 **efdcfsid**

**Convert floating-point double-precision from signed integer doubleword**

**efdcfsid** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtI64ToFP64}(rB_{0:63}, \text{SIGN})$$

The signed integer doubleword in **rB** 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** 

Scalar DPFP APU	User
-----------------	------

 **efdcfuf**

**Convert floating-point double-precision from unsigned fraction**

**efdcfuf** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtI32ToFP64}(rB_{32:63}, \text{UNSIGN}, F)$$

The unsigned fractional low element in **rB** is converted to a double-precision floating-point value using the current rounding mode and the result is placed into **rD**.


Exceptions:

None.

**efdcfui** Scalar DPFP APU User **efdcfui**

**Convert floating-point double-precision from unsigned integer**

**efdcfui** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtSI32ToFP64}(rB_{32:63}, \text{UNSIGN}, I)$$

The unsigned integer low element in **rB** is converted to a double-precision floating-point value using the current rounding mode and the result is placed into **rD**.


Exceptions:

None.

**efdcfuid** Scalar DPFP APU User **efdcfuid**

**Convert floating-point double-precision from unsigned integer doubleword**

**efdcfuid** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtI64ToFP64}(rB_{0:63}, \text{UNSIGN})$$

The unsigned integer doubleword in **rB** 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** Scalar DPFP APU User **efdcmpeq**

**Floating-point double-precision compare equal**

**efdcmpeq** **crfD,rA,rB**



$$aI \leftarrow rA_{0:63}$$

$$bI \leftarrow rB_{0:63}$$

$$\text{if } (aI = bI) \text{ then } cI \leftarrow 1$$


else  $cl \leftarrow 0$   
 $CR_{4*crD:4*crD+3} \leftarrow \text{undefined} || cl || \text{undefined} || \text{undefined}$

**rA** is compared against **rB**. If **rA** is equal to **rB**, the bit in the **crfD** is set, otherwise it is cleared. Comparison ignores the sign of 0 (+0 = -0).

Exceptions:

If the contents of **rA** or **rB** 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** Scalar DFPF APU User **efdcmpgt**  
**Floating-point double-precision compare greater than**




$a1 \leftarrow rA_{0:63}$   
 $b1 \leftarrow rB_{0:63}$   
 if ( $a1 > b1$ ) then  $cl \leftarrow 1$   
 else  $cl \leftarrow 0$   
 $CR_{4*crD:4*crD+3} \leftarrow \text{undefined} || cl || \text{undefined} || \text{undefined}$

**rA** is compared against **rB**. If **rA** is greater than **rB**, the bit in the **crfD** is set, otherwise it is cleared. Comparison ignores the sign of 0 (+0 = -0).

Exceptions:

If the contents of **rA** or **rB** 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** **efdcmplt**  
**Floating-point double-precision compare less than**



$a1 \leftarrow rA_{0:63}$   
 $b1 \leftarrow rB_{0:63}$   
 if ( $a1 < b1$ ) then  $cl \leftarrow 1$   
 else  $cl \leftarrow 0$   
 $CR_{4*crD:4*crD+3} \leftarrow \text{undefined} || cl || \text{undefined} || \text{undefined}$

**rA** is compared against **rB**. If **rA** is less than **rB**, the bit in the **crfD** is set, otherwise it is cleared. Comparison ignores the sign of 0 (+0 = -0).


Exceptions:

If the contents of **rA** or **rB** 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** Scalar DPFP APU User **efdctsf**

**Convert floating-point double-precision to signed fraction**

**efdctsf** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{SIGN}, \text{ROUND}, \text{F})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DPFP APU User **efdctsi**

**Convert floating-point double-precision to signed integer**

**efdctsi** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{SIGN}, \text{ROUND}, \text{I})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DPFP APU User **efdctsidz**

**Convert floating-point double-precision to signed integer doubleword with round toward zero**

**efdctsidz** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtFP64ToI64Sat}(rB_{0:63}, \text{SIGN}, \text{TRUNC})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DPFP APU User **efdctsiz**

**Convert floating-point double-precision to signed integer with round toward zero**

**efdctsiz** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{SIGN}, \text{TRUNC}, \text{I})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DPFP APU User **efdctuf**

**Convert floating-point double-precision to unsigned fraction**

**efdctuf** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{UNSIGN}, \text{ROUND}, \text{F})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DPFP APU User **efdctui**

**Convert floating-point double-precision to unsigned integer**

**efdctui** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{UNSIGN}, \text{ROUND}, \text{I})$$

The double-precision floating-point value in **rB** 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 **rB** 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** Scalar DFP APU User **efdctuidz**

**Convert floating-point double-precision to unsigned integer doubleword with round toward zero**

**efdctuidz** **rD,rB**



$$rD_{0:63} \leftarrow \text{CnvtFP64ToI64Sat}(rB_{0:63}, \text{UNSIGN}, \text{TRUNC})$$

The double-precision floating-point value in **rB** 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 **rB** 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.

**efdctui** Scalar DFP APU User **efdctui**

**Convert floating-point double-precision to unsigned integer with round toward zero**

**efdctui** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP64ToI32Sat}(rB_{0:63}, \text{UNSIGN}, \text{TRUNC}, \text{I})$$

The double-precision floating-point value in **rB** 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 **rB** 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**

Scalar DPFP APU	User
-----------------	------

**efddiv**

**Floating-point double-precision divide**

**efddiv** **rD,rA,rB**



$$rD_{0:63} \leftarrow rA_{0:63} \div_{dp} rB_{0:63}$$

**rA** is divided by **rB** and the result is stored in **rD**. If **rB** is a NaN or infinity, the result is a properly signed zero. Otherwise, if **rB** is a zero (or a denormalized number optionally transformed to zero by the implementation), or if **rA** is either NaN or infinity, the result is either *pmax* ( $a_{sign}==b_{sign}$ ), or *nmax* ( $a_{sign}!=b_{sign}$ ). 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 **rD**.

Exceptions:

If the contents of **rA** or **rB** are Infinity, Denorm, or NaN, or if both **rA** and **rB** 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 **rB** is +/-0 and the content of **rA** 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**

Scalar DPFP APU	User
-----------------	------

**efdmul**

**Floating-point double-precision multiply**

**efdmul** **rD,rA,rB**



0 0 0 1 0 0	rD	rA	rB	0 1 0 1 1 1 0 1 0 0 0
-------------	----	----	----	-----------------------

$$rD_{0:63} \leftarrow rA_{0:63} \times_{dp} rB_{0:63}$$

rA is multiplied by rB and the result is stored in rD. If rA or rB are zero (or a denormalized number optionally transformed to zero by the implementation), the result is a properly signed zero. Otherwise, if rA or rB are either NaN or infinity, the result is either *pmax* ( $a_{sign}=b_{sign}$ ), or *nmax* ( $a_{sign} \neq b_{sign}$ ). 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 rD.

Exceptions:

If the contents of rA or rB 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**

Scalar DPDFP APU	User
------------------	------

**efdnabs**

**Floating-point double-precision negative absolute value**

**efdnabs**

rD,rA

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD	rA	0	0	0	0	0	0	1	0	1	1	1	0	0	1	0	1								

$$rD_{0:63} \leftarrow 0b1 \parallel rA_{1:63}$$

The sign bit of rA is set to 1 and the result is placed into rD.

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 rA 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**

Scalar DPFP APU	User
-----------------	------

**efdneg**

**Floating-point double-precision negate**

**efdneg** **rD,rA**



$$rD_{0:63} \leftarrow \neg rA_0 \parallel rA_{1:63}$$

The sign bit of **rA** is complemented and the result is placed into **rD**.

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 **rA** 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.


**efdsb**

Scalar DPFP APU	User
-----------------	------

**efdsb**

**Floating-point double-precision subtract**

**efdsb** **rD,rA,rB**



$$rD_{0:63} \leftarrow rA_{0:63} -_{dp} rB_{0:63}$$

**rB** is subtracted from **rA** and the result is stored in **rD**. If **rA** is NaN or infinity, the result is either *pmax* ( $a_{sign}=0$ ), or *nmax* ( $a_{sign}=1$ ). Otherwise, If **rB** is NaN or infinity, the result is either *nmax* ( $b_{sign}=0$ ), or *pmax* ( $b_{sign}=1$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in **rD**. 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 **rA** or **rB** 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.

**efdsteq** Scalar DPFP APU User **efdsteq**  
**Floating-point double-precision test equal**

**efdsteq** **crfD,rA,rB**



```

aI ← rA0:63
bI ← rB0:63
if (aI = bI) then cI ← 1
else cI ← 0
CR4*crD:4*crD+3 ← undefined || cI || undefined || undefined
    
```


**rA** is compared against **rB**. If **rA** is equal to **rB**, 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 generated during the execution of **efdsteq** If strict IEEE 754 compliance is required, the program should use **efdcmpaq**.

Implementation note: In an implementation, the execution of **efdsteq** is likely to be faster than the execution of **efdcmpaq**.

**efdstgt** Scalar DPFP APU User **efdstgt**  
**Floating-point double-precision test greater than**

**efdstgt** **crfD,rA,rB**



```

aI ← rA0:63
bI ← rB0:63
if (aI > bI) then cI ← 1
else cI ← 0
CR4*crD:4*crD+3 ← undefined || cI || undefined || undefined
    
```

**rA** is compared against **rB**. If **rA** is greater than **rB**, 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 generated during the execution of **efdstgt**. If strict IEEE 754 compliance is required, the program should use **efdcmpgt**.



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

**efdtstlt**

Scalar DPFP APU	User
-----------------	------

**efdtstlt**  
**Floating-point double-precision test less than**  
**efdtstlt** **crfD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	<b>crfD</b>		0	0	<b>rA</b>				<b>rB</b>				0	1	0	1	1	1	1	1	1	1	0	1					

```

a1 ← rA0:63
b1 ← rB0:63
if (a1 < b1) then c1 ← 1
else c1 ← 0
CR4*crD:4*crD+3 ← undefined || c1 || undefined || undefined
    
```

rA is compared against rB. If rA is less than rB, 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 generated during the execution of **efdtstlt**. If strict IEEE 754 compliance is required, the program should use **efdcmpflt**.

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

**efsabs**

Scalar SPFP APU	User
-----------------	------

**efsabs**  
**Floating-Point Absolute Value**  
**efsabs** **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	<b>rD</b>				<b>rA</b>				0	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1	0	0

```

rD32:63 ← 0b0 || rA33:63
    
```

The sign bit of rA is cleared and the result is placed into rD.

It is implementation dependent if invalid values for rA (NaN, Denorm, Infinity) are detected and exceptions are taken.

**efsadd**

Scalar SPFP APU	User
-----------------	------

**efsadd**  
**Floating-Point Add**  
**efsadd** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	<b>rD</b>				<b>rA</b>				<b>rB</b>				0	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0

```

rD32:63 ← rA32:63 +sp rB32:63
    
```



The single-precision floating-point value of **rA** is added to **rB** and the result is stored in **rD**.

If an overflow condition is detected or the contents of **rA** or **rB** 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 **rA** or **rB** 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** Scalar SPFP APU User **efscfsf**  
**Convert Floating-Point from Signed Fraction**

**efscfsf** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{F})$$


The signed fractional value in **rB** 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** Scalar SPFP APU User **efscfsi**  
**Convert Floating-Point from Signed Integer**

**efscfsi** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtSI32ToFP32Sat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{I})$$


The signed integer value in **rB** 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

**efscfuf** Scalar SPFP APU User **efscfuf**  
**Convert Floating-Point from Unsigned Fraction**

**efscfuf** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{F})$$

The unsigned fractional value in **rB** 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** Scalar SPFP APU User **efscfui**  
**Convert Floating-Point from Unsigned Integer**

**efscfui** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{I})$$


The unsigned integer value in **rB** 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** Scalar SPFP APU User **efscmpeq**  
**Floating-Point Compare Equal**

**efscmpeq** **crD,rA,rB**



$$aI \leftarrow rA_{32:63}$$

$$bI \leftarrow rB_{32:63}$$

$$\text{if } (aI = bI) \text{ then } cI \leftarrow 1$$

$$\text{else } cI \leftarrow 0$$

$$CR_{4*crD:4*crD+3} \leftarrow \text{undefined} \parallel cI \parallel \text{undefined} \parallel \text{undefined}$$

The value in **rA** is compared against **rB**. If **rA** equals **rB**, the **crD** 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 **rA** or **rB** are +infinity, -infinity, denorm or NaN

**efscmpgt** Scalar SPFP APU User **efscmpgt**  
**Floating-Point Compare Greater Than**

**efscmpgt** **crD,rA,rB**



$$aI \leftarrow rA_{32:63}$$

$$bI \leftarrow rB_{32:63}$$

$$\text{if } (aI > bI) \text{ then } cI \leftarrow 1$$

else  $cl \leftarrow 0$   
 $CR_{4*crD:4*crD+3} \leftarrow \text{undefined} || cl || \text{undefined} || \text{undefined}$

The value in **rA** is compared against **rB**. If **rA** is greater than **rB**, the bit in the **crD** 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 **rA** or **rB** are +infinity, -infinity, denorm or NaN


**efscmplt**

Scalar SPFP APU	User
-----------------	------

**efscmplt**

**Floating-Point Compare Less Than**

**efscmplt** **crD,rA,rB**



$a1 \leftarrow rA_{32:63}$   
 $b1 \leftarrow rB_{32:63}$   
 if ( $a1 < b1$ ) then  $cl \leftarrow 1$   
 else  $cl \leftarrow 0$   
 $CR_{4*crD:4*crD+3} \leftarrow \text{undefined} || cl || \text{undefined} || \text{undefined}$

The value in **rA** is compared against **rB**. If **rA** is less than **rB**, the bit in the **crD** 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 **rA** or **rB** are +infinity, -infinity, denorm or NaN


**efscfsf**

Scalar SPFP APU	User
-----------------	------

**efscfsf**

**Convert Floating-Point to Signed Fraction**

**efscfsf** **rD,rB**



$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{ROUND}, \text{F})$

The single-precision floating-point value in **rB** 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 **rB** are +infinity., -infinity, denorm, or NaN, or **rB** cannot be represented in the target format
- FINXS, FG, FX if the result is inexact

**efsctsi** Scalar SPFP APU User **efsctsi**

**Convert Floating-Point to Signed Integer**

**efsctsi** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{ROUND}, \text{I})$$

The single-precision floating-point value in **rB** 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 **rB** are +infinity, -infinity, denorm, or NaN, or **rB** cannot be represented in the target format
- FINXS, FG, FX if the result is inexact

**efsctsiz** Scalar SPFP APU User **efsctsiz**

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

**efsctsiz** **rD,rB**




$$rD_{32-63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{TRUNC}, \text{I})$$

The single-precision floating-point value in **rB** 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.

**efsctuf** Scalar SPFP APU User **efsctuf**

**Convert Floating-Point to Unsigned Fraction**

**efsctuf** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{ROUND}, \text{F})$$

The single-precision floating-point value in **rB** 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 **rB** are +infinity, -infinity, denorm, or NaN, or **rB** cannot be represented in the target format
- FINXS, FG, FX if the result is inexact

**efsctui** Scalar SPFP APU User **efsctui**

**Convert Floating-Point to Unsigned Integer**

**efsctui** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{ROUND}, I)$$

The single-precision floating-point value in **rB** 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 **rB** are +infinity, -infinity, denorm, or NaN, or **rB** cannot be represented in the target format
- FINXS, FG, FX if the result is inexact

**efsctuiZ** Scalar SPFP APU User **efsctuiZ**

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

**efsctuiZ** **rD,rB**



$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{TRUNC}, I)$$


The single-precision floating-point value in **rB** 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 **rB** are +infinity, -infinity, denorm, or NaN, or **rB** cannot be represented in the target format
- FINXS, FG, FX if the result is inexact

**efsddiv** Scalar SPFP APU User **efsddiv**  
**Floating-Point Divide**

**efsddiv** **rD,rA,rB**



$$rD_{32:63} \leftarrow rA_{32:63} \div_{sp} rB_{32:63}$$


The single-precision floating-point value in **rA** is divided by **rB** and the result is stored in **rD**.  
 If an overflow is detected, or **rB** is a denorm (or 0 value), or **rA** is a NaN or Infinity and **rB** 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 **rA** or **rB** 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** Scalar SPFP APU User **efsmul**  
**Floating-Point Multiply**

**efsmul** **rD,rA,rB**



$$rD_{32:63} \leftarrow rA_{32:63} \times_{sp} rB_{32:63}$$

The single-precision floating-point value in **rA** is multiplied by **rB** and the result is stored in **rD**.  
 If an overflow is detected the result is an appropriately signed maximum floating-point value.  
 If one of **rA** or **rB** 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 **rA** or **rB** 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 **rA** or **rB** 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

**efsnabs**

Scalar SPFP APU	User
-----------------	------

**efsnabs**

**Floating-Point Negative Absolute Value**

**efsnabs**

**rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						rA						0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1	0	1

$$rD_{32:63} \leftarrow 0b1 \parallel rA_{33:63}$$

The sign bit of **rA** is set and the result is stored in **rD**. It is implementation dependent if invalid values for **rA** (NaN, Denorm, Infinity) are detected and exceptions are taken.

**efsneg**

Scalar SPFP APU	User
-----------------	------

**efsneg**

**Floating-Point Negate**

**efsneg**

**rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD						rA						0	0	0	0	0	0	0	0	0	1	0	1	1	0	0	0	1	1	0

$$rD_{32:63} \leftarrow \neg rA_{32} \parallel rA_{33:63}$$

The sign bit of **rA** is complemented and the result is stored in **rD**. It is implementation dependent if invalid values for **rA** (NaN, Denorm, Infinity) are detected and exceptions are taken.

**efssub**

Scalar SPFP APU	User
-----------------	------

**efssub**

**Floating-Point Subtract**

**efssub**

**rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31						
0	0	0	1	0	0	rD						rA						rB						0	1	0	1	1	0	0	0	0	0	0	0	0	1

$$rD_{32:63} \leftarrow rA_{32:63} -_{sp} rB_{32:63}$$

The single-precision floating-point value in **rB** is subtracted from that in **rA** and the result is stored in **rD**.

If an overflow condition is detected or the contents of **rA** or **rB** 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 **rA** or **rB** 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** Scalar SPFP APU User **efststeq**

**Floating-Point Test Equal**

**efststeq** **crD,rA,rB**



```

al ← rA32:63
bl ← rB32:63
if (al = bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← undefined || cl || undefined || undefined
    
```


The value in **rA** is compared against **rB**. If **rA** equals **rB**, 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 execution of **efststeq**. If strict IEEE 754 compliance is required, the program should use **efscmpeq**.

**efststgt** Scalar SPFP APU User **efststgt**

**Floating-Point Test Greater Than**

**efststgt** **crD,rA,rB**



```

al ← rA32:63
bl ← rB32:63
if (al > bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← undefined || cl || undefined || undefined
    
```


If **rA** is greater than **rB**, 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 **efscmpgt**.

**efstslt** Scalar SPFP APU User **efstslt**

**Floating-Point Test Less Than**

**efstslt** **crD,rA,rB**



```

al ← rA32:63
bl ← rB32:63
    
```



if (a1 < b1) then c1 ← 1  
 else c1 ← 0  
 CR<sub>4\*crD:4\*crD+3</sub> ← undefined || c1 || undefined || undefined

If **rA** is less than **rB**, the bit in the **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 **efststltt**. If strict IEEE 754 compliance is required, the program should use **efscmplt**.


**eqv**

Book E	User
--------	------

**eqv**

**Equivalent**

**eqv** **rA,rS,rB(Rc=0)**  
**eqv.** **rA,rS,rB(Rc=1)**



result<sub>0:63</sub> ← rS ≡ rB  
 if Rc=1 then do  
     LT ← result<sub>32:63</sub> < 0  
     GT ← result<sub>32:63</sub> > 0  
     EQ ← result<sub>32:63</sub> = 0  
     CR0 ← LT || GT || EQ || SO  
 rA ← result

The contents of **rS** are XORed with the contents of **rB** and the one's complement of the result is placed into **rA**.

Other registers altered:

- CR0 (if Rc=1)


**evabs**

SPE APU	User
---------	------

**evabs**

**Vector absolute value**

**evabs** **rD,rA**




rD<sub>0:31</sub> ← ABS(rA<sub>0:31</sub>)  
 rD<sub>32:63</sub> ← ABS(rA<sub>32:63</sub>)

The absolute value of each element of **rA** is placed in the corresponding elements of **rD**. An absolute value of 0x8000\_0000 (most negative number) returns 0x8000\_0000. No overflow is detected.



Figure 87. Vector absolute value (evabs)



evaddiw


SPE APU	User
---------	------

evaddiw

Vector add immediate word

evaddiw

rD,rB,UIMM

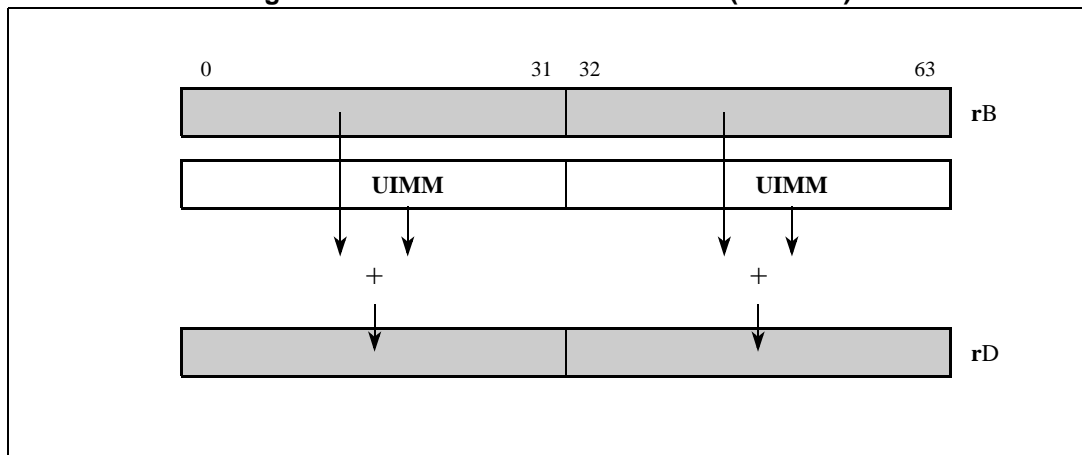


$$rD_{0:31} \leftarrow rB_{0:31} + \text{EXTZ}(UIMM) \quad // \text{ Modulo sum}$$

$$rD_{32:63} \leftarrow rB_{32:63} + \text{EXTZ}(UIMM) \quad // \text{ Modulo sum}$$

UIMM is zero-extended and added to both the high and low elements of rB and the results are placed in rD. Note that the same value is added to both elements of the register. UIMM is 5 bits.

Figure 88. Vector add immediate word (evaddiw)



**evaddsmiaaw** SPE APU User **evaddsmiaaw**

**Vector add signed, modulo, integer to accumulator word**

**evaddsmiaaw** **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD						rA						0	0	0	0	0	1	0	0	1	1	0	0	1	0	0	1

$$rD_{0:31} \leftarrow ACC_{0:31} + rA_{0:31}$$


$$rD_{32:63} \leftarrow ACC_{32:63} + rA_{32:63}$$

$$ACC_{0:63} \leftarrow rD_{0:63}$$

Each word element in **rA** is added to the corresponding element in the accumulator and the results are placed in **rD** and into the accumulator.

Other registers altered: ACC

**Figure 89. Vector add signed, modulo, integer to accumulator word (evaddsmiaaw)**



**evaddssiaaw** SPE APU User **evaddssiaaw**

**Vector add signed, saturate, integer to accumulator word**

**evaddssiaaw** **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD						rA						0	0	0	0	0	1	0	0	1	1	0	0	0	0	0	1

```
// high
temp0:63 ← EXTS(ACC0:31) + EXTS(rA0:31)
ovh ← temp31 ⊕ temp32
rD0:31 ← SATURATE(ovh, temp31, 0x80000000, 0x7ffffff, temp32:63)
```


```
// low
temp0:63 ← EXTS(ACC32:63) + EXTS(rA32:63)
ovl ← temp31 ⊕ temp32
rD32:63 ← SATURATE(ovl, temp31, 0x80000000, 0x7ffffff, temp32:63)
```

$ACC_{0:63} \leftarrow rD_{0:63}$   
 $SPEFSCR_{OVH} \leftarrow ovh$   
 $SPEFSCR_{OV} \leftarrow ovl$   
 $SPEFSCR_{SOVH} \leftarrow SPEFSCR_{SOVH} | ovh$   
 $SPEFSCR_{SOV} \leftarrow SPEFSCR_{SOV} | ovl$

Each signed integer word element in rA 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 rD and the accumulator. Any overflow or underflow is recorded in the SPEFSCR overflow and summary overflow bits.

Other registers altered: SPEFSCR ACC

**Figure 90. Vector add signed, saturate, integer to accumulator word (evaddssiaaw)**



**evaddumiaaw**

SPE	APU	User
-----	-----	------

**evaddumiaaw**

**Vector add unsigned, modulo, integer to accumulator word**

**evaddumiaaw**      rD,rA


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	1	0	0	1	1	0	0	1	0	0	0

$rD_{0:31} \leftarrow ACC_{0:31} + rA_{0:31}$   
 $rD_{32:63} \leftarrow ACC_{32:63} + rA_{32:63}$   
 $ACC_{0:63} \leftarrow rD_{0:63}$

Each unsigned integer word element in rA is added to the corresponding element in the accumulator and the results are placed in rD and the accumulator.

Other registers altered: ACC

Figure 91. Vector add unsigned, modulo, integer to accumulator word (evaddumiaaw)



evaddusiaaw

SPE APU	User
---------	------

evaddusiaaw

Vector add unsigned, saturate, integer to accumulator word

evaddusiaaw

rD,rA

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	1	0	0	1	1	0	0	0	0	0	0

```

// high
temp0:63 ← EXTZ(ACC0:31) + EXTZ(rA0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, temp31, 0xffffffff, 0xffffffff, temp32:63)

// low
temp0:63 ← EXTZ(ACC32:63) + EXTZ(rA32:63)
ovl ← temp31
rD32:63 ← SATURATE(ovl, temp31, 0xffffffff, 0xffffffff, temp32:63)


ACC0:63 ← rD0:63

SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

Each unsigned integer word element in rA 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. Any overflow is recorded in the SPEFSCR overflow and summary overflow bits.

Other registers altered: SPEFSCR ACC

**Figure 92. Vector add unsigned, saturate, integer to accumulator word (evaddusiaaw)**



**evaddw**

SPE APU	User
---------	------

**evaddw**

**Vector add word**

**evaddw**

**rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31							
0	0	0	1	0	0	rD										rA						rB				0	1	0	0	0	0	0	0	0	0	0	0	0

$$rD_{0:31} \leftarrow rA_{0:31} + rB_{0:31}$$


// Modulo sum

$$rD_{32:63} \leftarrow rA_{32:63} + rB_{32:63}$$

// Modulo sum

The corresponding elements of rA and rB are added and the results are placed in rD. The sum is a modulo sum.


**Figure 93. Vector add word (evaddw)**



**evand** SPE APU User **evand**

**Vector AND**

**evand** **rD,rA,rB**




$$rD_{0:31} \leftarrow rA_{0:31} \& rB_{0:31} \quad // \text{ Bitwise AND}$$

$$rD_{32:63} \leftarrow rA_{32:63} \& rB_{32:63} \quad // \text{ Bitwise AND}$$

The corresponding elements of **rA** and **rB** are ANDed bitwise and the results are placed in the corresponding element of **rD**.


**Figure 94. Vector AND (evand)**



**evandc** SPE APU User **evandc**

**Vector AND with complement**

**evandc** **rD,rA,rB**

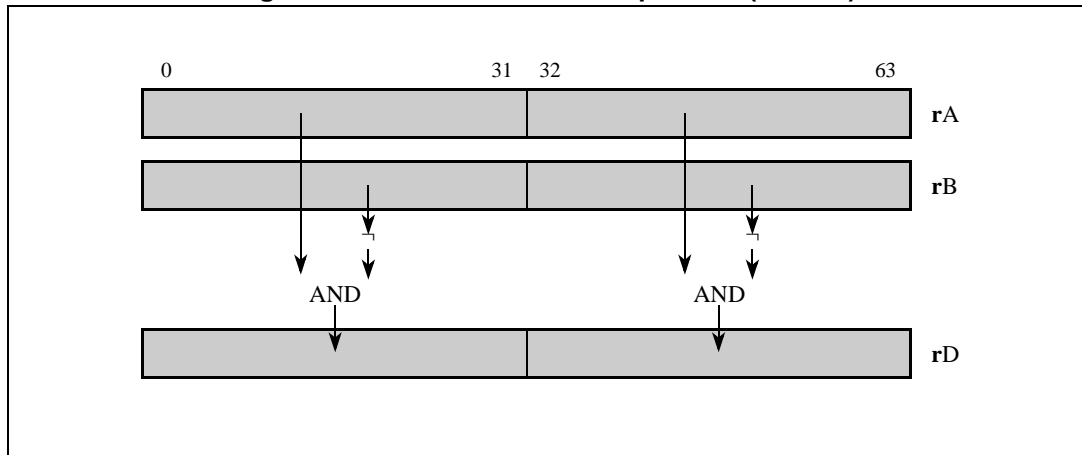


$$rD_{0:31} \leftarrow rA_{0:31} \& (\neg rB_{0:31}) \quad // \text{ Bitwise ANDC}$$

$$rD_{32:63} \leftarrow rA_{32:63} \& (\neg rB_{32:63}) \quad // \text{ Bitwise ANDC}$$

The word elements of **rA** and are ANDed bitwise with the complement of the corresponding elements of **rB**. The results are placed in the corresponding element of **rD**.

Figure 95. Vector AND with complement (evandc)



**evcmpeq**

SPE APU	User
---------	------

**evcmpeq**

Vector compare equal  
**evcmpeq**

**crD,rA,rB**

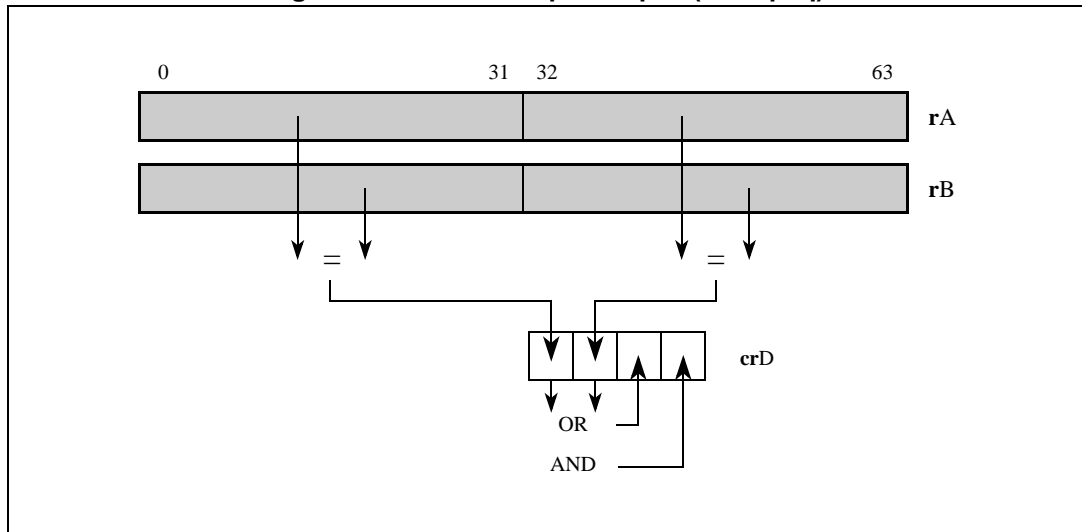
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	<b>crD</b>		0	0	<b>rA</b>				<b>rB</b>				0	1	0	0	0	1	1	0	1	0	0				

```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah = bh) then ch ← 1
else ch ← 0
if (al = bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

The most significant bit in **crD** is set if the high-order element of **rA** is equal to the high-order element of **rB**; it is cleared otherwise. The next bit in **crD** is set if the low-order element of **rA** is equal to the low-order element of **rB** and cleared otherwise. The last two bits of **crD** are set to the OR and AND of the result of the compare of the high and low elements.

Figure 96. Vector Compare Equal (evcmpgts)



evcmpgts

SPE APU	User
---------	------

evcmpgts

Vector compare greater than signed

evcmpgts

crD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	crD		0	0	rA						rB						0	1	0	0	0	1	1	0	0	0	1

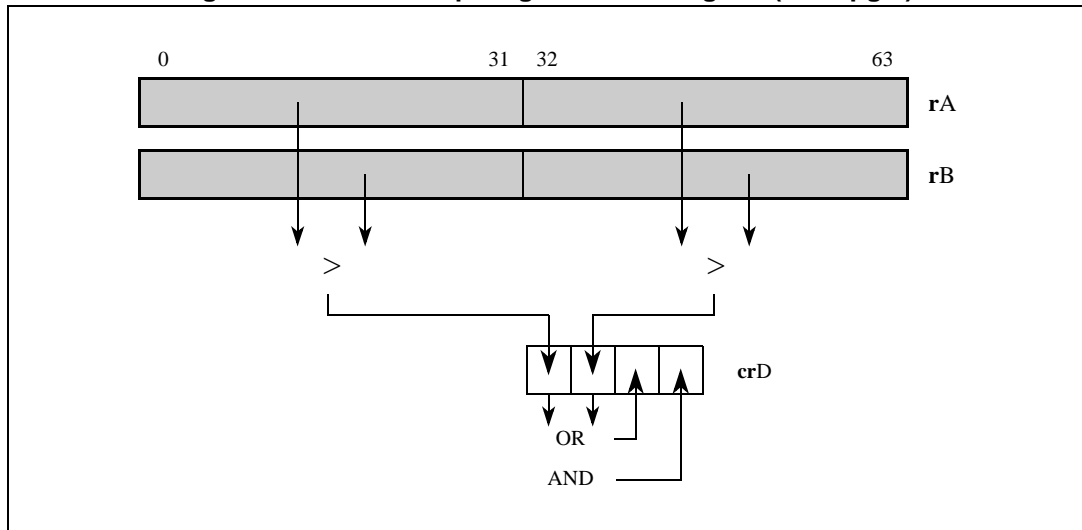
```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah > bh) then ch ← 1
else ch ← 0
if (al > bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

The most significant bit in **crD** is set if the high-order element of **rA** is greater than the high-order element of **rB**; it is cleared otherwise. The next bit in **crD** is set if the low-order element of **rA** is greater than the low-order element of **rB** and cleared otherwise. The last two bits of **crD** are set to the OR and AND of the result of the compare of the high and low elements.



Figure 97. Vector compare greater than signed (evcmpgts)



**evcmpgtu** SPE APU User **evcmpgtu**

**Vector compare greater than unsigned**

**evcmpgtu** **crD,rA,rB**

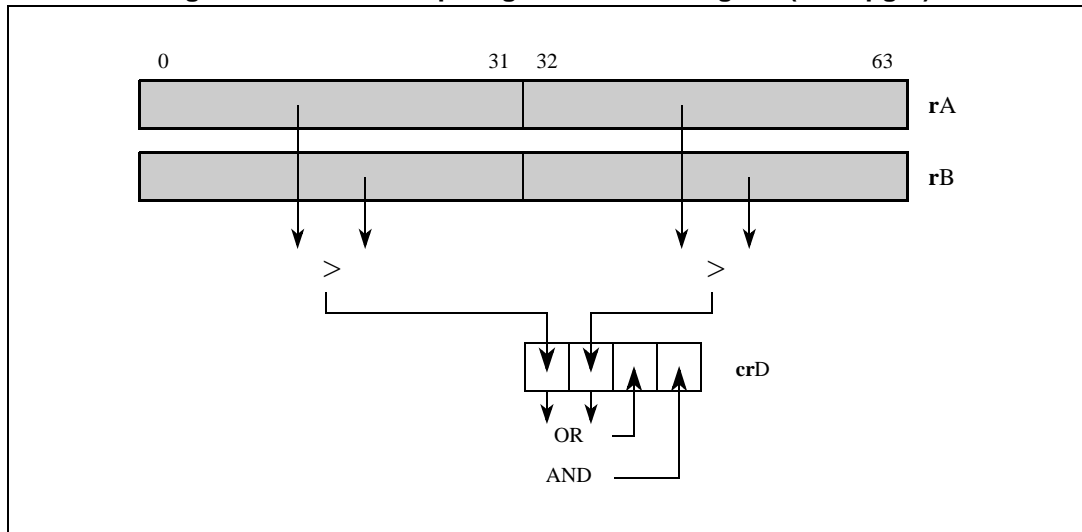
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	crD		0	0	rA						rB						0	1	0	0	0	1	1	0	0	0	0	0

```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah >U bh) then ch ← 1
else ch ← 0
if (al >U bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

The most significant bit in **crD** is set if the high-order element of **rA** is greater than the high-order element of **rB**; it is cleared otherwise. The next bit in **crD** is set if the low-order element of **rA** is greater than the low-order element of **rB** and cleared otherwise. The last two bits of **crD** are set to the OR and AND of the result of the compare of the high and low elements.

Figure 98. Vector compare greater than unsigned (evcmpgtu)



evcmplt

SPE APU	User
---------	------

evcmplt

Vector compare less than signed

evcmplt

crD,rA,rB

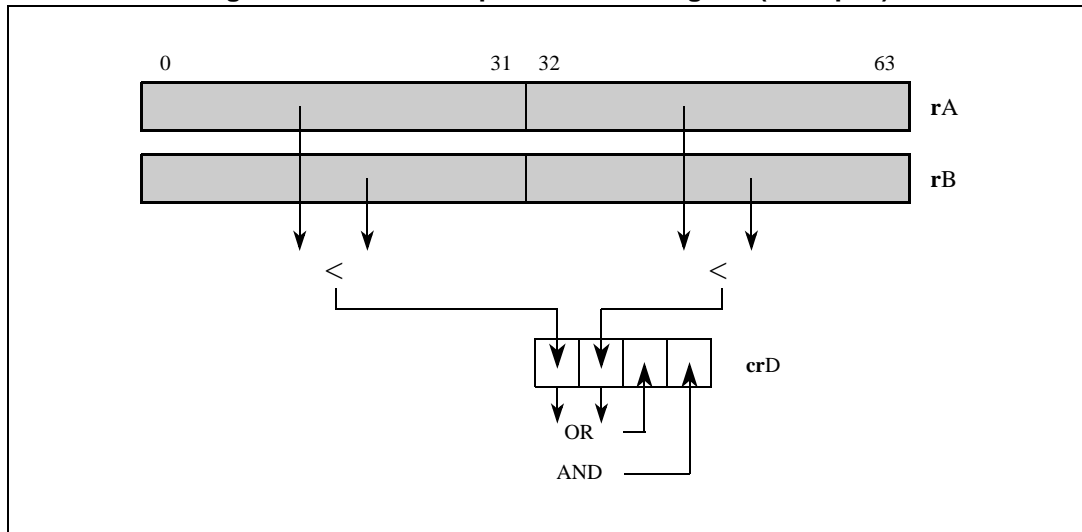
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	crD		0	0	rA						rB						0	1	0	0	0	1	1	0	0	1	1

```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah < bh) then ch ← 1
else ch ← 0
if (al < bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

The most significant bit in **crD** is set if the high-order element of **rA** is less than the high-order element of **rB**; it is cleared otherwise. The next bit in **crD** is set if the low-order element of **rA** is less than the low-order element of **rB** and cleared otherwise. The last two bits of **crD** are set to the OR and AND of the result of the compare of the high and low elements.

Figure 99. Vector compare less than signed (evcmpltu)



evcmpltu

SPE APU	User
---------	------

evcmpltu

Vector compare less than unsigned

evcmpltu

crD,rA,rB

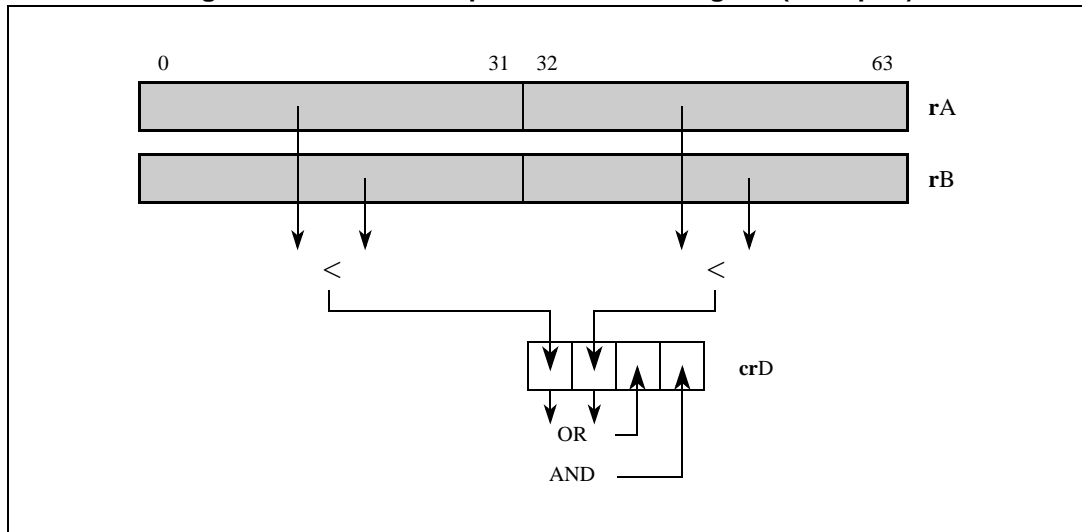
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	crD		0	0	rA				rB				0	1	0	0	0	0	1	1	0	0	1	0	0	0	0

```


ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah <U bh) then ch ← 1
else ch ← 0
if (al <U bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

The most significant bit in **crD** is set if the high-order element of **rA** is less than the high-order element of **rB**; it is cleared otherwise. The next bit in **crD** is set if the low-order element of **rA** is less than the low-order element of **rB** and cleared otherwise. The last two bits of **crD** are set to the OR and AND of the result of the compare of the high and low elements.

Figure 100. Vector compare less than unsigned (evcmpltu)



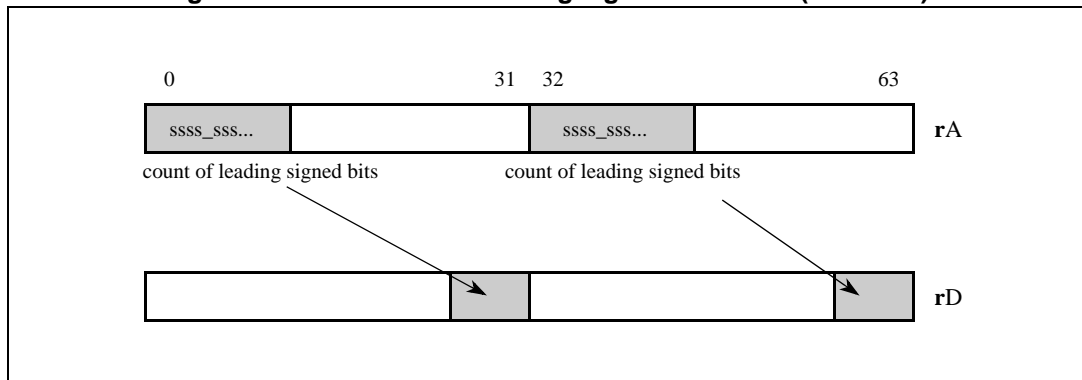
**evcntlsw** SPE APU User **evcntlsw**  
**Vector count leading signed bits word**  
**evcntlsw** **rD,rA**



The leading sign bits in each element of rA are counted, and the respective count is placed into each element of rD.

**evcntlzw** is used for unsigned operands; **evcntlsw** is used for signed operands.

Figure 101. Vector count leading signed bits word (evcntlsw)




**evcntlzw**

SPE APU	User
---------	------

**evcntlzw**


**Vector count leading zeros word**

**evcntlzw**                      **rD,rA**



The leading zero bits in each element of **rA** are counted, and the respective count is placed into each element of **rD**.

**Figure 102. Vector count leading zeros word (evcntlzw)**




**evdivws**

SPE APU	User
---------	------

**evdivws**

**Vector divide word signed**

**evdivws**                      **rD,rA,rB**



```


dividendh ← rA0:31
dividendl ← rA32:63
divisorh ← rB0:31
divisorl ← rB32:63
rD0:31 ← dividendh ÷ divisorh
rD32:63 ← dividendl ÷ divisorl
ovh ← 0
ovl ← 0
if ((dividendh < 0) & (divisorh = 0)) then
    rD0:31 ← 0x80000000
    ovh ← 1
else if ((dividendh >= 0) & (divisorh = 0)) then
    rD0:31 ← 0x7FFFFFFF
    ovh ← 1
else if ((dividendh = 0x80000000) & (divisorh = 0xFFFF_FFFF)) then
    rD0:31 ← 0x7FFFFFFF
    ovh ← 1
    
```

```

if ((dividendh < 0) & (divisorl = 0)) then
    rD32:63 ← 0x80000000
    ovl ← 1
else if ((dividendh >= 0) & (divisorl = 0)) then
    rD32:63 ← 0x7FFFFFFF
    ovl ← 1
else if ((dividendh = 0x80000000) & (divisorl = 0xFFFF_FFFF)) then
    rD32:63 ← 0x7FFFFFFF
    ovl ← 1
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The two dividends are the two elements of the contents of rA. The two divisors are the two elements of the contents of rB. The resulting two 32-bit quotients on each element are placed into rD. The 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 103. Vector divide word signed (evdivws)



evdivwu

SPE APU	User
---------	------

evdivwu

Vector divide word unsigned

evdivwu                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																											

```


dividendh ← rA0:31
dividendl ← rA32:63
divisorh ← rB0:31
divisorl ← rB32:63
rD0:31 ← dividendh ÷ divisorh
rD32:63 ← dividendl ÷ divisorl
ovh ← 0
ovl ← 0
if (divisorh = 0) then
    
```

```

rD0:31 = 0xFFFFFFFF
ovh ← 1
if (divisorl = 0) then
    rD32:63 ← 0xFFFFFFFF
    ovl ← 1
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The two dividends are the two elements of the contents of **rA**. The two divisors are the two elements of the contents of **rB**. 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 104. Vector divide word unsigned (evdivwu)**



**eveqv** SPE APU User **eveqv**  
**Vector equivalent**  
**eveqv** **rD,rA,rB**

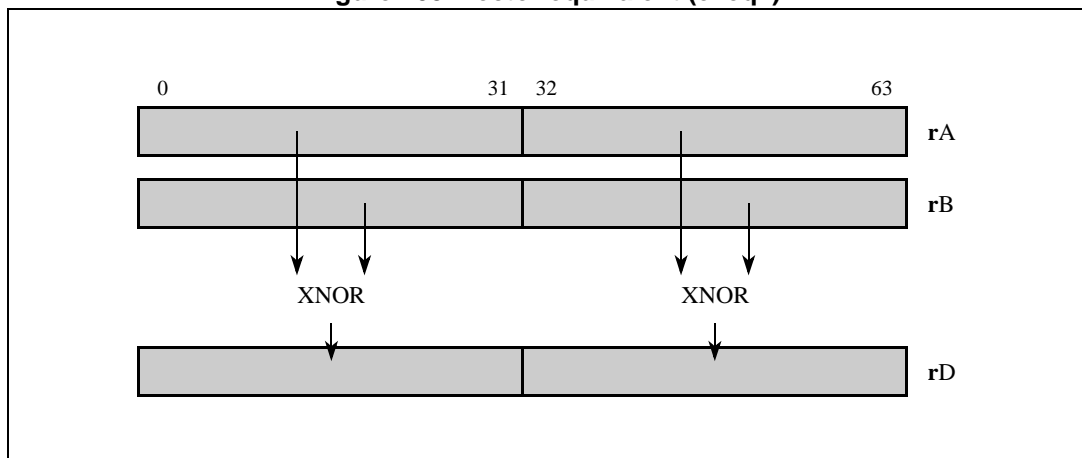
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	0	0	1	1	0	0	1

$$rD_{0:31} \leftarrow rA_{0:31} \oplus rB_{0:31} \quad // \text{ Bitwise XNOR}$$

$$rD_{32:63} \leftarrow rA_{32:63} \oplus rB_{32:63} \quad // \text{ Bitwise XNOR}$$

The corresponding elements of rA & rB are XNORed bitwise, & the results are placed in rD.

**Figure 105. Vector equivalent (eveqv)**



**evextsb** SPE APU User **evextsb**  
**Vector extend sign byte**  
**evextsb** **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						rA						0	0	0	0	0	0	0	0	1	0	0	0	0	0	1	0	1	0


$$rD_{0:31} \leftarrow \text{EXTS}(rA_{24:31})$$

$$rD_{32:63} \leftarrow \text{EXTS}(rA_{56:63})$$

The signs of the byte in each of the elements in rA are extended, and the results are placed in rD.



Figure 106. Vector extend sign byte (evextsb)



evextsh

SPE APU	User
---------	------

evextsh

Vector extend sign half word

evextsh

rD,rA


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	0	0	1	0	0	0	0	0	1	0	1	1

$$rD_{0:31} \leftarrow \text{EXTS}(rA_{16:31})$$

$$rD_{32:63} \leftarrow \text{EXTS}(rA_{48:63})$$

The signs of the half words in each of the elements in rA are extended, and the results are placed in rD.

Figure 107. Vector extend sign half word (evextsh)



**evfsabs**


Vector SPFP APU	User
-----------------	------

**evfsabs**

**Vector floating-point single-precision absolute value**

**evfsabs**

**rD,rA**



$$rD_{0:31} \leftarrow 0b0 \parallel rA_{1:31}$$

$$rD_{32:63} \leftarrow 0b0 \parallel rA_{33:63}$$

The sign bit of each element in **rA** is set to 0 and the results are placed into **rD**.

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**


Vector SPFP APU	User
-----------------	------

**evfsadd**

**Vector floating-point single-precision add**

**evfsadd**

**rD,rA,rB**



$$rD_{0:31} \leftarrow rA_{0:31} +_{sp} rB_{0:31}$$

$$rD_{32:63} \leftarrow rA_{32:63} +_{sp} rB_{32:63}$$

Each single-precision floating-point element of **rA** is added to the corresponding element of **rB** and the results are stored in **rD**. If an element of **rA** is NaN or infinity, the corresponding result is either *pmax* ( $a_{sign}=0$ ), or *nmax* ( $a_{sign}=1$ ). Otherwise, if an element of **rB** is NaN or infinity, the corresponding result is either *pmax* ( $b_{sign}=0$ ), or *nmax* ( $b_{sign}=1$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in the corresponding element of **rD**. If an underflow occurs, +0 (for rounding modes RN, RZ, RP) or -0 (for rounding mode RM) is stored in the corresponding element of **rD**.

Exceptions:

If the contents of either element of **rA** or **rB** 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).

**evfscfsf** Vector SPFP APU User **evfscfsf**  
**Vector convert floating-point single-precision from signed fraction**  
**evfscfsf** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{0:31}, \text{SIGN}, \text{UPPER}, \text{F})$$


$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{F})$$

Each signed fractional element of rB is converted to a single-precision floating-point value using the current rounding mode and the results are placed into the corresponding elements of rD.

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.

**evfscfsi** Vector SPFP APU User **evfscfsi**  
**Vector convert floating-point single-precision from signed integer**  
**evfscfsi** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtSI32ToFP32Sat}(rB_{0:31}, \text{SIGN}, \text{UPPER}, \text{I})$$

$$rD_{32:63} \leftarrow \text{CnvtSI32ToFP32Sat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{I})$$

Each signed integer element of rB 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 rD.




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** Vector SPFP APU User **evfscfuf**

**Vector convert floating-point single-precision from unsigned fraction**

**evfscfuf** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{0:31}, \text{UNSIGN}, \text{UPPER}, \text{F})$$

$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{F})$$

Each unsigned fractional element of **rB** is converted to a single-precision floating-point value using the current rounding mode and the results are placed into the corresponding elements of **rD**.


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.

**evfscfui** Vector SPFP APU User **evfscfui**

**Vector convert floating-point single-precision from unsigned integer**

**evfscfui** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{0:31}, \text{UNSIGN}, \text{UPPER}, \text{I})$$

$$rD_{32:63} \leftarrow \text{CnvtI32ToFP32Sat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{I})$$

Each unsigned integer element of **rB** 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 **rD**.

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 SPFP APU User **evfscmpeq**

**Vector floating-point single-precision compare equal**

**evfscmpeq** **crfD,rA,rB**



```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah = bh) then ch ← 1
else ch ← 0
if (al = bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared against the corresponding element of **rB**. If **rA** equals **rB**, the **crfD** bit is set, otherwise it is cleared. Comparison ignores the sign of 0 (+0 = -0).


Exceptions:

If the contents of either element of **rA** or **rB** 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** Vector SPFP APU User **evfscmpgt**

**Vector floating-point single-precision compare greater than**

**evfscmpgt** **crfD,rA,rB**



```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah > bh) then ch ← 1
else ch ← 0
if (al > bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared against the corresponding element of **rB**. If **rA** is greater than **rB**, 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 **rA** or **rB** 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** Vector SPFP APU User **evfscmplt**  
**Vector floating-point single-precision compare less than**  
**evfscmplt** **crfD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	<b>crfD</b>		0	0	<b>rA</b>				<b>rB</b>				0 1 0 1 0 0 0 1 1 0 1													

```

ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah < bh) then ch ← 1
else ch ← 0
if (al < bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared against the corresponding element of **rB**. If **rA** is less than **rB**, 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 **rA** or **rB** 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.



**evfscfsf** Vector SPFP APU User **evfscfsf**  
**Vector convert floating-point single-precision to signed fraction**

**evfscfsf** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{SIGN}, \text{UPPER}, \text{ROUND}, \text{F})$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{ROUND}, \text{F})$$

Each single-precision floating-point element in **rB** 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 **rB** 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.

**evfscfsi** Vector SPFP APU User **evfscfsi**  
**Vector convert floating-point single-precision to signed integer**

**evfscfsi** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{SIGN}, \text{UPPER}, \text{ROUND}, \text{I})$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{ROUND}, \text{I})$$

Each single-precision floating-point element in **rB** 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 **rB** 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.


**evfsctsiz**

Vector SPFP APU	User
-----------------	------

**evfsctsiz**

**Vector convert floating-point single-precision to signed integer with round toward zero**

**evfsctsiz** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{SIGN}, \text{UPPER}, \text{TRUNC}, \text{I})$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{SIGN}, \text{LOWER}, \text{TRUNC}, \text{I})$$

Each single-precision floating-point element in **rB** 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 **rB** 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**

Vector SPFP APU	User
-----------------	------

**evfsctuf**

**Vector convert floating-point single-precision to unsigned fraction**

**evfsctuf** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{UNSIGN}, \text{UPPER}, \text{ROUND}, \text{F})$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{ROUND}, \text{F})$$



Each single-precision floating-point element in **rB** 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 **rB** 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.

**evfsctui** Vector SPFP APU User **evfsctui**

**Vector convert floating-point single-precision to unsigned integer**

**evfsctui** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{UNSIGN}, \text{UPPER}, \text{ROUND}, I)$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{ROUND}, I)$$

Each single-precision floating-point element in **rB** 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 **rB** 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** Vector SPFP APU User **evfsctuiZ**

**Vector convert floating-point single-precision to unsigned integer with round toward zero**

**evfsctuiZ** **rD,rB**



$$rD_{0:31} \leftarrow \text{CnvtFP32ToISat}(rB_{0:31}, \text{UNSIGN}, \text{UPPER}, \text{TRUNC}, I)$$

$$rD_{32:63} \leftarrow \text{CnvtFP32ToISat}(rB_{32:63}, \text{UNSIGN}, \text{LOWER}, \text{TRUNC}, I)$$

Each single-precision floating-point element in **rB** 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 **rB** 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.

**evfdiv**

Vector SPFP APU	User
-----------------	------

**evfdiv**

**Vector floating-point single-precision divide**

**evfdiv**

**rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						0	1	0	1	0	0	0	1	0	0	1

$$rD_{0:31} \leftarrow rA_{0:31} \div_{sp} rB_{0:31}$$

$$rD_{32:63} \leftarrow rA_{32:63} \div_{sp} rB_{32:63}$$

Each single-precision floating-point element of **rA** is divided by the corresponding element of **rB** and the result is stored in **rD**. If an element of **rB** is a NaN or infinity, the corresponding result is a properly signed zero. Otherwise, if an element of **rB** is a zero (or a denormalized number optionally transformed to zero by the implementation), or if an element of **rA** is either NaN or infinity, the corresponding result is either *pmax* ( $a_{sign}=b_{sign}$ ), or *nmax* ( $a_{sign} \neq b_{sign}$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in the corresponding element of **rD**. If an underflow occurs, +0 or -0 (as appropriate) is stored in the corresponding element of **rD**.

Exceptions:

If the contents of **rA** or **rB** are Infinity, Denorm, or NaN, or if both **rA** and **rB** 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 **rB** is  $\pm 0$  and the content of **rA** 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** Vector SPFP APU User **evfsmul**

**Vector floating-point single-precision multiply**

**evfsmul** **rD,rA,rB**



$$rD_{0:31} \leftarrow rA_{0:31} \times_{sp} rB_{0:31}$$

$$rD_{32:63} \leftarrow rA_{32:63} \times_{sp} rB_{32:63}$$

Each single-precision floating-point element of **rA** is multiplied with the corresponding element of **rB** and the result is stored in **rD**. If an element of **rA** or **rB** 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 **rA** or **rB** are either NaN or infinity, the corresponding result is either *pmax* ( $a_{sign}=b_{sign}$ ), or *nmax* ( $a_{sign} \neq b_{sign}$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in the corresponding element of **rD**. If an underflow occurs, +0 or -0 (as appropriate) is stored in the corresponding element of **rD**.

Exceptions:

If the contents of either element of **rA** or **rB** 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**


Vector SPFP APU	User
-----------------	------

**evfsnabs**

**Vector floating-point single-precision negative absolute value**

**evfsnabs**

**rD,rA**



$$rD_{0:31} \leftarrow 0b1 \parallel rA_{1:31}$$

$$rD_{32:63} \leftarrow 0b1 \parallel rA_{33:63}$$

The sign bit of each element in rA is set to 1 and the results are placed into rD.

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.

**evfsneg**


Vector SPFP APU	User
-----------------	------

**evfsneg**

**Vector floating-point single-precision negate**

**evfsneg**

**rD,rA**



$$rD_{0:31} \leftarrow \neg rA_0 \parallel rA_{1:31}$$

$$rD_{32:63} \leftarrow \neg rA_{32} \parallel rA_{33:63}$$

The sign bit of each element in rA is complemented and the results are placed into rD.

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** Vector SPFP APU User **evfssub**

**Vector floating-point single-precision subtract**

**evfssub** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31								
0	0	0	1	0	0	<b>rD</b>						<b>rA</b>						<b>rB</b>						0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	1

$$rD_{0:31} \leftarrow rA_{0:31} -_{sp} rB_{0:31}$$

$$rD_{32:63} \leftarrow rA_{32:63} -_{sp} rB_{32:63}$$

Each single-precision floating-point element of **rB** is subtracted from the corresponding element of **rA** and the results are stored in **rD**. If an element of **rA** is NaN or infinity, the corresponding result is either *pmax* ( $a_{sign}=0$ ), or *nmax* ( $a_{sign}=1$ ). Otherwise, if an element of **rB** is NaN or infinity, the corresponding result is either *nmax* ( $b_{sign}=0$ ), or *pmax* ( $b_{sign}=1$ ). Otherwise, if an overflow occurs, *pmax* or *nmax* (as appropriate) is stored in the corresponding element of **rD**. If an underflow occurs, +0 (for rounding modes RN, RZ, RP) or -0 (for rounding mode RM) is stored in the corresponding element of **rD**.

Exceptions:

If the contents of either element of **rA** or **rB** 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 SPFP APU User **evfststeq**

**Vector floating-point single-precision test equal**

**evfststeq** **crfD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	<b>crfD</b>		0	0	<b>rA</b>						<b>rB</b>						0	1	0	1	0	0	1	1	1	1	0

$$ah \leftarrow rA_{0:31}$$

$$al \leftarrow rA_{32:63}$$



```

bh ← rB0:31
bl ← rB32:63
if (ah = bh) then ch ← 1
else ch ← 0
if (al = bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared against the corresponding element of **rB**. If **rA** equals **rB**, 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 **evfststgq**. If strict IEEE 754 compliance is required, the program should use **evfscmpeq**.

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

**evfststgt** Vector SPFP APU User **evfststgt**

**Vector floating-point single-precision test greater than**

**evfststgt** **crfD,rA,rB**



```


ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah > bh) then ch ← 1
else ch ← 0
if (al > bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared against the corresponding element of **rB**. If **rA** is greater than **rB**, 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** Vector SPFP APU User **evfststlt**  
**Vector floating-point single-precision test less than**  
**evfststlt** **crfD,rA,rB**



```


ah ← rA0:31
al ← rA32:63
bh ← rB0:31
bl ← rB32:63
if (ah < bh) then ch ← 1
else ch ← 0
if (al < bl) then cl ← 1
else cl ← 0
CR4*crD:4*crD+3 ← ch || cl || (ch | cl) || (ch & cl)
    
```

Each element of **rA** is compared with the corresponding element of **rB**. If **rA** is less than **rB**, 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**.

**evldd** SPE, Vector SPFP, Scalar DFP APU User **evldd**  
**Vector load double word into double word**  
**evldd** **rD,d(rA)**



1. **d** = UIMM \* 8

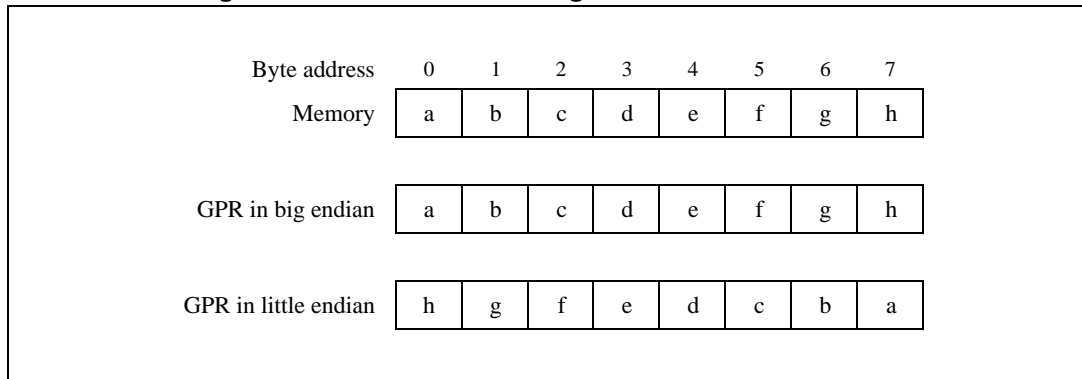
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*8)
rD ← MEM(EA, 8)
    
```

The double word addressed by EA is loaded from memory and placed in rD.


*Figure 108* shows how bytes are loaded into rD as determined by the endian mode.

**Figure 108. evladd results in big- and little-endian modes**



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

**evladdx** SPE, Vector SPFP, Scalar DPFP APUs User **evladdx**  
**Vector load double word into double word indexed**  
**evladdx** **rD,rA,rB**

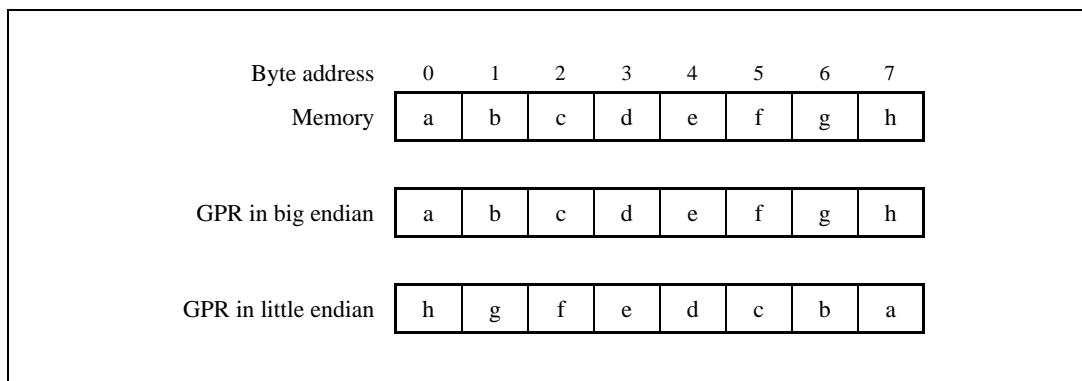


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 rD ← MEM(EA, 8)

The double word addressed by EA is loaded from memory and placed in rD.

Figure 109 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 109. evladdx results in big- and little-endian modes**



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



**evldh**

SPE APU	User
---------	------

**evldh**

**Vector load double into four half words**

**evldh** **rD,d(rA)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD						rA						UIMM <sup>(1)</sup>				0	1	1	0	0	0	0	0	0	1	0	1

1.  $d = UIMM * 8$

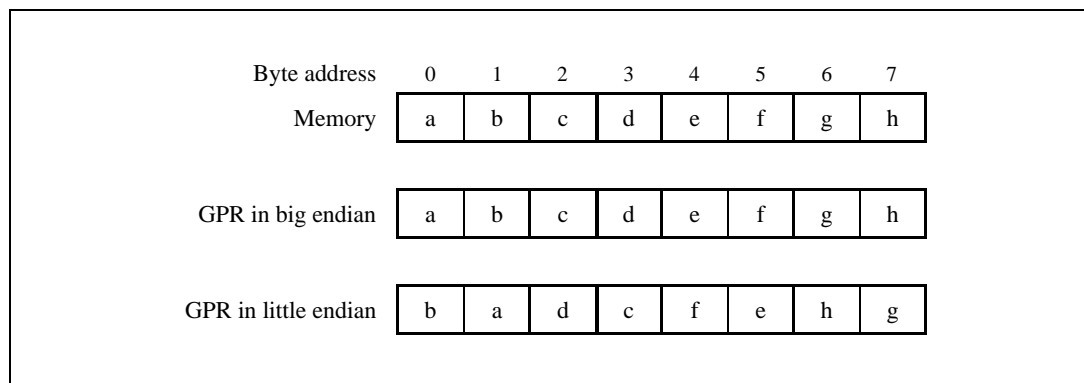
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*8)
rD0:15 ← MEM(EA, 2)
rD16:31 ← MEM(EA+2,2)
rD32:47 ← MEM(EA+4,2)
rD48:63 ← MEM(EA+6,2)
    
```

The double word addressed by EA is loaded from memory and placed in rD.

The [Figure 110](#) shows how bytes are loaded into rD as determined by the endian mode.

**Figure 110. evldh Results in Big- and Little-Endian Modes**



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

**evldhx**

SPE APU	User
---------	------

**evldhx**

**Vector Load Double into Four Half Words Indexed**

**evldhx** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD						rA						rB				0	1	1	0	0	0	0	0	0	1	0	0

```

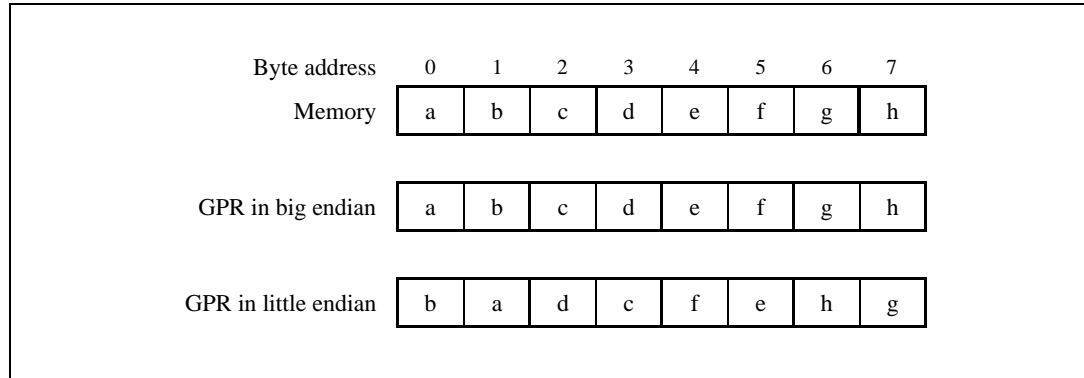
if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:15 ← MEM(EA, 2)
rD16:31 ← MEM(EA+2,2)
rD32:47 ← MEM(EA+4,2)
rD48:63 ← MEM(EA+6,2)
    
```



The double word addressed by EA is loaded from memory and placed in rD.


Figure 111 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 111. evldhx results in big- and little-endian modes**



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

**evldw** SPE APU User **evldw**  
**Vector load double into two words**  
**evldw** **rD,d(rA)**



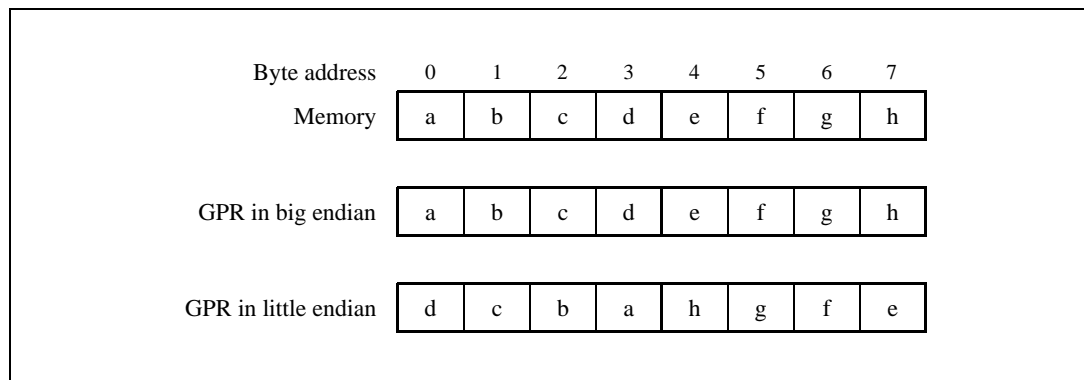
1. **d** = UIMM \* 8

if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*8)  
 rD<sub>0:31</sub> ← MEM(EA, 4)  
 rD<sub>32:63</sub> ← MEM(EA+4, 4)

The double word addressed by EA is loaded from memory and placed in rD.

Figure 112 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 112. evldw results in big- and little-endian modes**



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


**evldwx**

SPE APU	User
---------	------

**evldwx**

**Vector load double into two words indexed**

**evldwx** **rD,rA,rB**



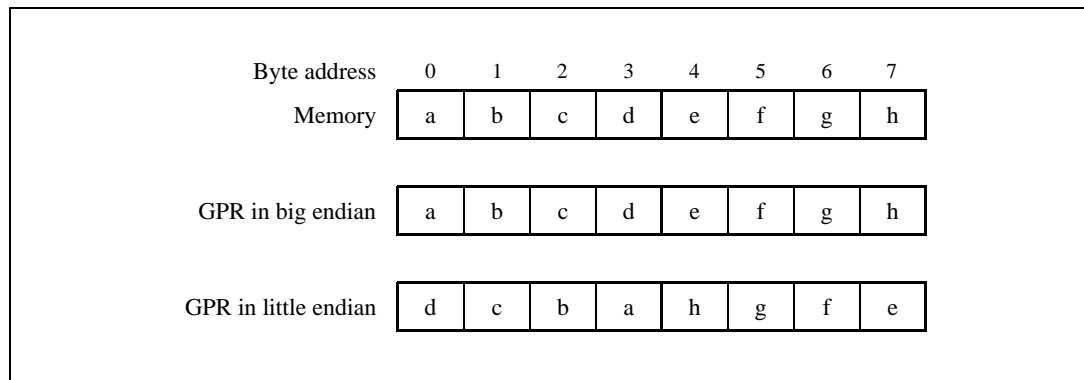
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:31 ← MEM(EA, 4)
rD32:63 ← MEM(EA+4, 4)
    
```

The double word addressed by EA is loaded from memory and placed in rD.

Figure 113 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 113. evldwx results in big- and little-endian modes**



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


**evlhhesplat**

SPE APU	User
---------	------

**evlhhesplat**

**Vector load half word into half words even and splat**

**evlhhesplat** **rD,d(rA)**



1. **d** = UIMM \* 2


```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*2)
rD0:15 ← MEM(EA,2)
rD16:31 ← 0x0000
rD32:47 ← MEM(EA,2)
rD48:63 ← 0x0000
    
```

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

Figure 114 shows how bytes are loaded into rD as determined by the endian mode.

Figure 114. evlhhesplat results in big- and little-endian modes



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

**evlhhesplatx**

SPE APU	User
---------	------

**evlhhesplatx**  
**Vector load half word into half words even and splat indexed**  
**evlhhesplatx**      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						0	1	1	0	0	0	0	1	0	0	0


```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:15 ← MEM(EA,2)
rD16:31 ← 0x0000
rD32:47 ← MEM(EA,2)
rD48:63 ← 0x0000
    
```

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

Figure 115 shows how bytes are loaded into rD as determined by the endian mode.

Figure 115. evlhhesplatx results in big- and little-endian modes




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

**evlhossplat** SPE APU User **evlhossplat**

**Vector load half word into half word odd signed and splat**

**evlhossplat** **rD,d(rA)**




1. **d** = UIMM \* 2

if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*2)  
 rD<sub>0:31</sub> ← EXTS(MEM(EA,2))  
 rD<sub>32:63</sub> ← 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 116 shows how bytes are loaded into rD as determined by the endian mode.


Figure 116. evlhossplat 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.

**evlhossplatx** SPE APU User **evlhossplatx**  
**Vector load half word into half word odd signed and splat indexed**  
**evlhossplatx** **rD,rA,rB**




```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:31 ← EXTS(MEM(EA,2))
rD32:63 ← 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 117 shows how bytes are loaded into rD as determined by the endian mode.


**Figure 117. evlhossplatx 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 APU User **evlhhousplat**  
**Vector load half word into half word odd unsigned and splat**  
**evlhhousplat** **rD,d(rA)**



1. d = UIMM \* 2

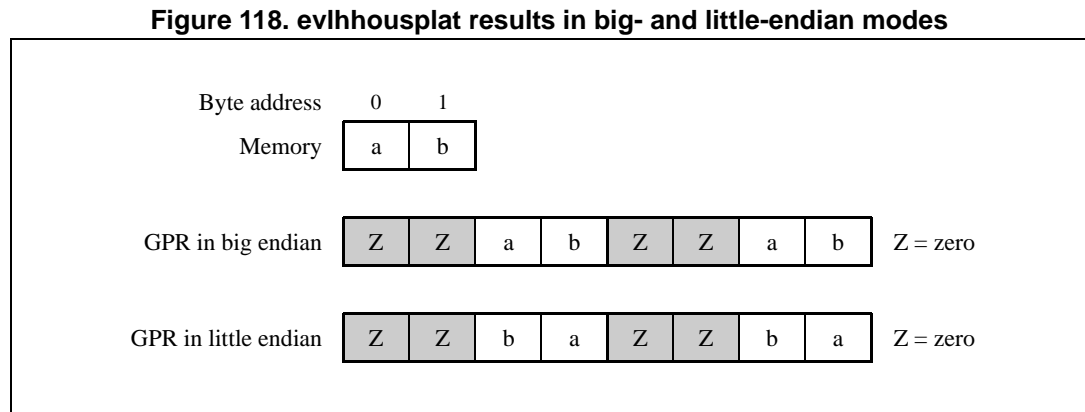
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*2)
rD0:15 ← 0x0000
    
```

$rD_{16:31} \leftarrow \text{MEM}(\text{EA},2)$   
 $rD_{32:47} \leftarrow 0x0000$   
 $rD_{48:63} \leftarrow \text{MEM}(\text{EA},2)$

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

Figure 118 shows how bytes are loaded into rD as determined by the endian mode.




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

**evlhhousplatx**      SPE APU   User      **evlhhousplatx**

**Vector load half word into half word odd unsigned and splat indexed**

**evlhhousplatx**      **rD,rA,rB**




if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 $rD_{0:15} \leftarrow 0x0000$   
 $rD_{16:31} \leftarrow \text{MEM}(\text{EA},2)$   
 $rD_{32:47} \leftarrow 0x0000$   
 $rD_{48:63} \leftarrow \text{MEM}(\text{EA},2)$

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


Figure 119 shows how bytes are loaded into rD as determined by the endian mode.

Figure 119. evlhhousplatx results in big- and little-endian modes



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

**evlwhe** SPE APU User **evlwhe**  
**Vector load word into two half words even**  
**evlwhe** **rD,d(rA)**




1. d = UIMM \* 4

if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*4)  
 rD<sub>0:15</sub> ← MEM(EA,2)  
 rD<sub>16:31</sub> ← 0x0000  
 rD<sub>32:47</sub> ← MEM(EA+2,2)  
 rD<sub>48:63</sub> ← 0x0000

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

Figure 120 shows how bytes are loaded into rD as determined by the endian mode.

Figure 120. evlwhe results in big- and little-endian modes





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

**evlwhex** SPE APU User **evlwhex**

**Vector load word into two half words even indexed**

**evlwhex** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					0	1	1	0	0	0	1	0	0	0	0


```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:15 ← MEM(EA,2)
rD16:31 ← 0x0000
rD32:47 ← MEM(EA+2,2)
rD48:63 ← 0x0000
    
```

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

Figure 121 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 121. evlwhex results in big- and little-endian modes**



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

**evlwhos** SPE APU User **evlwhos**

**Vector load word into two half words odd signed (with sign extension)**

**evlwhos** **rD,d(rA)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					UIMM <sup>(1)</sup>					0	1	1	0	0	0	1	0	1	1	1

1. d = UIMM \* 4

```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*4)
    
```


$$rD_{0:31} \leftarrow \text{EXTS}(\text{MEM}(\text{EA},2))$$

$$rD_{32:63} \leftarrow \text{EXTS}(\text{MEM}(\text{EA}+2,2))$$

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

Figure 122 shows how bytes are loaded into rD as determined by the endian mode.

Figure 122. 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 APU User **evlwhosx**

**Vector load word into two half words odd signed indexed (with sign extension)**

**evlwhosx** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					0	1	1	0	0	0	1	0	1	1	0


```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:31 ← EXTS(MEM(EA,2))
rD32:63 ← EXTS(MEM(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 rD.

Figure 123 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 123. 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 APU | User **evlwhou**

**Vector load word into two half words odd unsigned (zero-extended)**

**evlwhou** **rD,d(rA)**



1. **d** = UIMM \* 4


```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*4)
rD0:15 ← 0x0000
rD16:31 ← MEM(EA,2)
rD32:47 ← 0x0000
rD48:63 ← 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 rD.


Figure 124 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 124. evlwhou results in big- and little-endian modes**



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

**evlwhoux** SPE APU User **evlwhoux**  
**Vector load word into two half words odd unsigned indexed (zero-extended)**  
**evlwhoux** **rD,rA,rB**




```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:15 ← 0x0000
rD16:31 ← MEM(EA,2)
rD32:47 ← 0x0000
rD48:63 ← 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 rD.


Figure 125 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 125. evlwhoux results in big- and little-endian modes**



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

**evlwhsplat** SPE APU User **evlwhsplat**  
**Vector load word into two half words and splat**  
**evlwhsplat** **rD,d(rA)**



1. d = UIMM \* 4

```

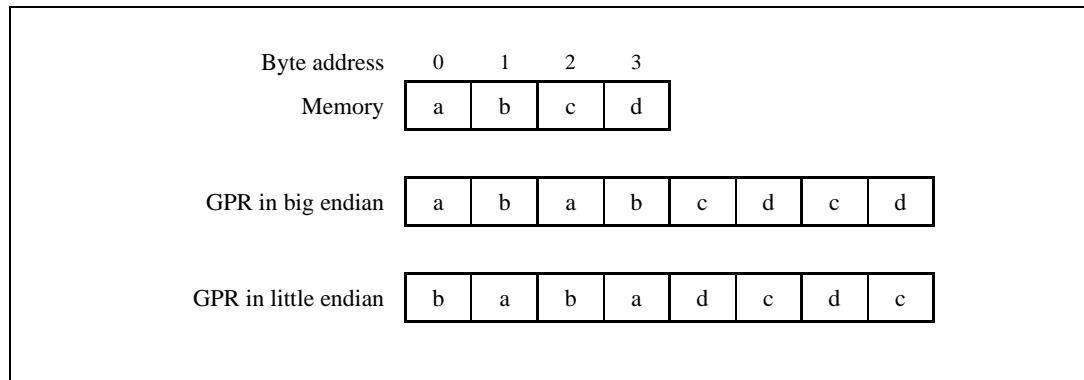
if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*4)
rD0:15 ← MEM(EA,2)
    
```

$rD_{16:31} \leftarrow \text{MEM}(\text{EA}, 2)$   
 $rD_{32:47} \leftarrow \text{MEM}(\text{EA}+2, 2)$   
 $rD_{48:63} \leftarrow \text{MEM}(\text{EA}+2, 2)$

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

Figure 126 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 126. evlwhsplat results in big- and little-endian modes**



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

**evlwhsplatx**


SPE APU	User
---------	------

**evlwhsplatx**

**Vector load word into two half words and splat indexed**

**evlwhsplatx**

**rD,rA,rB**

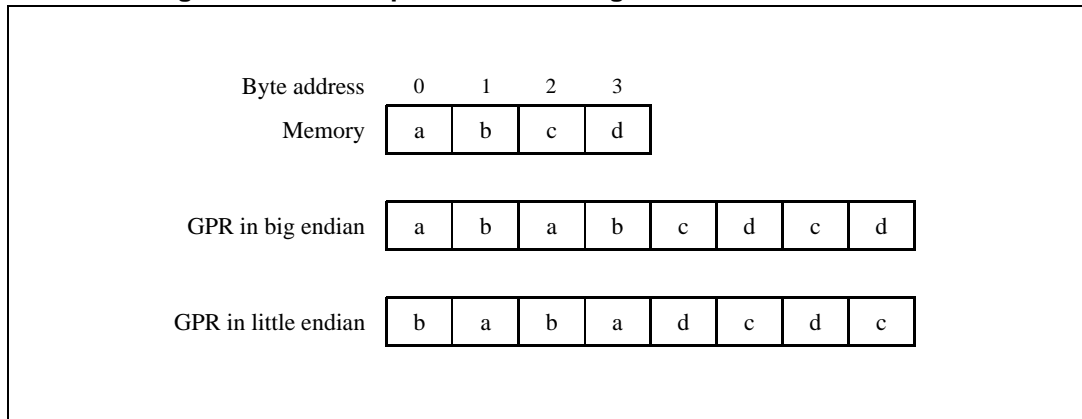


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 $rD_{0:15} \leftarrow \text{MEM}(\text{EA}, 2)$   
 $rD_{16:31} \leftarrow \text{MEM}(\text{EA}, 2)$   
 $rD_{32:47} \leftarrow \text{MEM}(\text{EA}+2, 2)$   
 $rD_{48:63} \leftarrow \text{MEM}(\text{EA}+2, 2)$

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

Figure 127 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 127. evlwhsplatx results in big- and little-endian modes**




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

**evlwwsplat** SPE APU User **evlwwsplat**

**Vector load word into word and splat**

**evlwwsplat** **rD,d(rA)**



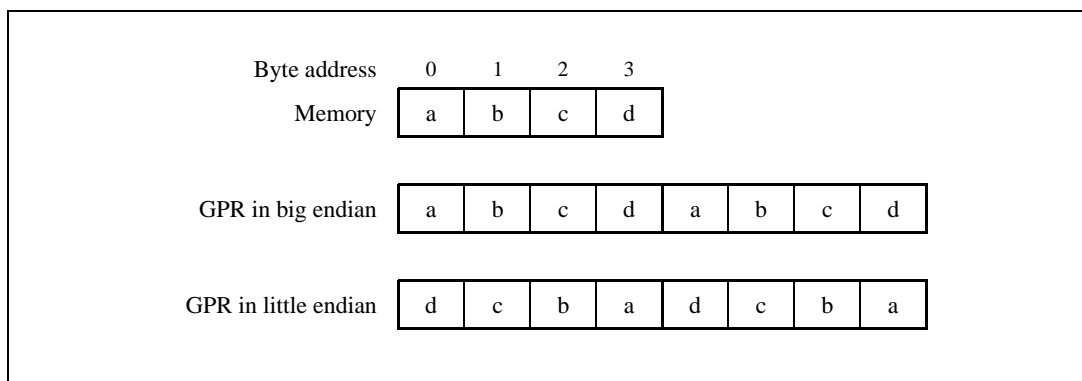
1.  $d = UIMM * 4$

if (rA = 0) then  $b \leftarrow 0$   
 else  $b \leftarrow (rA)$   
 $EA \leftarrow b + EXTZ(UIMM*4)$   
 $rD_{0:31} \leftarrow MEM(EA,4)$   
 $rD_{32:63} \leftarrow MEM(EA,4)$

The word addressed by EA is loaded from memory and placed in both elements of rD.

*Figure 128* shows how bytes are loaded into rD as determined by the endian mode.

**Figure 128. evlwwsplat results in big- and little-endian modes**



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


**evlwwsplatx**

SPE APU	User
---------	------

**evlwwsplatx**

Vector load word into word and splat indexed

**evlwwsplatx**                      **rD,rA,rB**



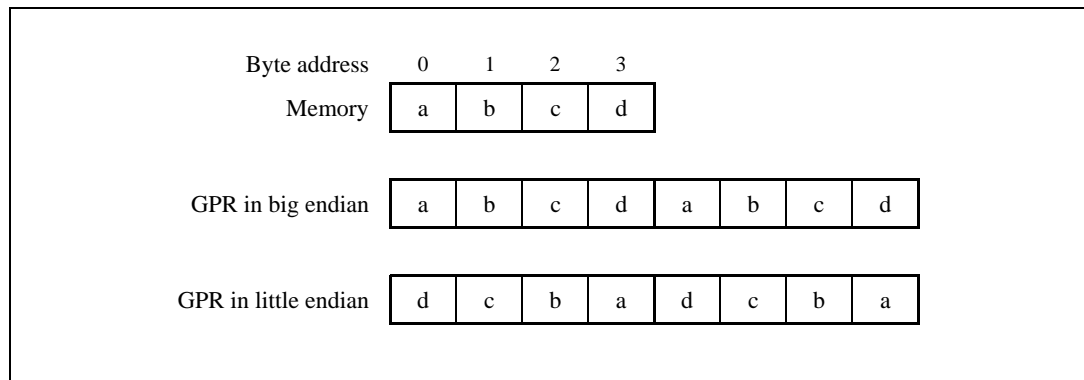
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
rD0:31 ← MEM(EA,4)
rD32:63 ← MEM(EA,4)
    
```

The word addressed by EA is loaded from memory and placed in both elements of rD.

Figure 129 shows how bytes are loaded into rD as determined by the endian mode.

**Figure 129. evlwwsplatx results in big- and little-endian modes**



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


**evmergehi**

SPE, Vector SPFP, scalar DPFP APUs	User
------------------------------------	------

**evmergehi**

Vector merge high

**evmergehi**                      **rD,rA,rB**




```

rD0:31 ← rA0:31
rD32:63 ← rB0:31
    
```

The high-order elements of rA and rB are merged and placed into rD, as shown in Figure 130.


Figure 130. High order element merging (evmergehi)



Note: A vector splat high can be performed by specifying the same register in rA and rB.

```

evmergehilo [SPE, Vector SPFP, Scalar DFPF APU] User evmergehilo
Vector merge high/low
evmergehilo rD,rA,rB
    
```




$$rD_{0:31} \leftarrow rA_{0:31}$$

$$rD_{32:63} \leftarrow rB_{32:63}$$

The high-order element of rA and the low-order element of rB are merged and placed into rD, as shown in Figure 131.

Figure 131. High order element merging (evmergehilo)



Application note: With appropriate specification of rA and rB, **evmergehi**, **evmergeho**, **evmergehilo**, and **evmergeholo** provide a full 32-bit permute of two source operands.



**evmergelo**


SPE APU	User
---------	------

**evmergelo**

**Vector merge low**

**evmergelo**

**rD,rA,rB**




$$rD_{0:31} \leftarrow rA_{32:63}$$

$$rD_{32:63} \leftarrow rB_{32:63}$$

The low-order elements of **rA** and **rB** are merged and placed in **rD**, as shown in [Figure 132](#).

**Figure 132. Low order element merging (evmergelo)**



*Note:* A vector splat low can be performed by specifying the same register in **rA** and **rB**.

**evmergelohi**


SPE APU	User
---------	------

**evmergelohi**

**Vector merge low/high**

**evmergelohi**

**rD,rA,rB**




$$rD_{0:31} \leftarrow rA_{32:63}$$

$$rD_{32:63} \leftarrow rB_{0:31}$$

The low-order element of **rA** and the high-order element of **rB** are merged and placed into **rD**, as shown in [Figure 133](#).

Figure 133. Low order element merging (evmergelohi)



Note: A vector swap can be performed by specifying the same register in rA and rB.

**evmhegsmfaa** SPE APU User **evmhegsmfaa**  
**Vector multiply half words, even, guarded, signed, modulo, fractional and accumulate**  
**evmhegsmfaa** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						1	0	1	0	0	1	0	1	0	1	1

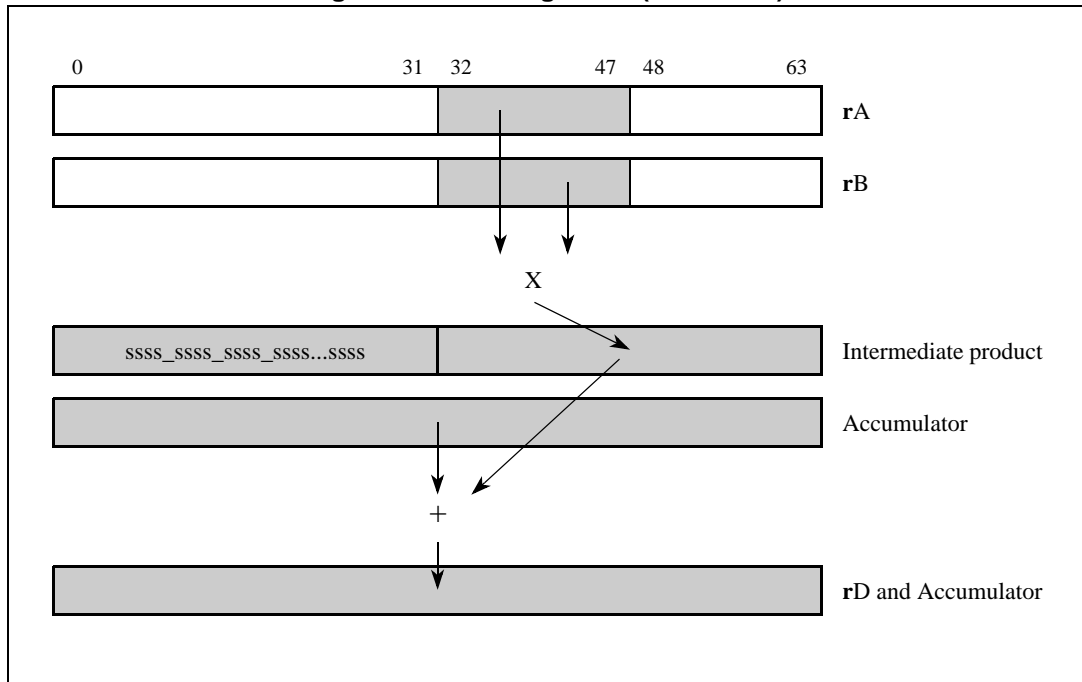
```
temp0:31 ← rA32:47 ×sf rB32:47
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 + temp0:63

// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered, half-word signed fractional elements in rA and rB are multiplied. The product is added to the contents of the 64-bit accumulator and the result is placed into rD and the accumulator.

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 134. evmhegsmfaa (even form)



evmhegsmfan

SPE APU	User
---------	------

evmhegsmfan

Vector multiply half words, even, guarded, signed, modulo, fractional and accumulate negative

evmhegsmfan                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0					rD					rA					rB		1	0	1	1	0	1	0	1	0	1	1


```
temp0:31 ← rA32:47 ×sf rB32:47
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered, half-word signed fractional elements in rA and rB are multiplied. The product is subtracted from the contents of the 64-bit accumulator and the result is placed into rD and the accumulator.

*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 135. evmhegsmfan (even form)



**evmhegsmiaa** SPE APU User **evmhegsmiaa**  
**Vector multiply half words, even, guarded, signed, modulo, integer and accumulate**  
**evmhegsmiaa** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31															
0	0	0	1	0	0	rD										rA										rB										1	0	1	0	0	1	0	1	0	0	1

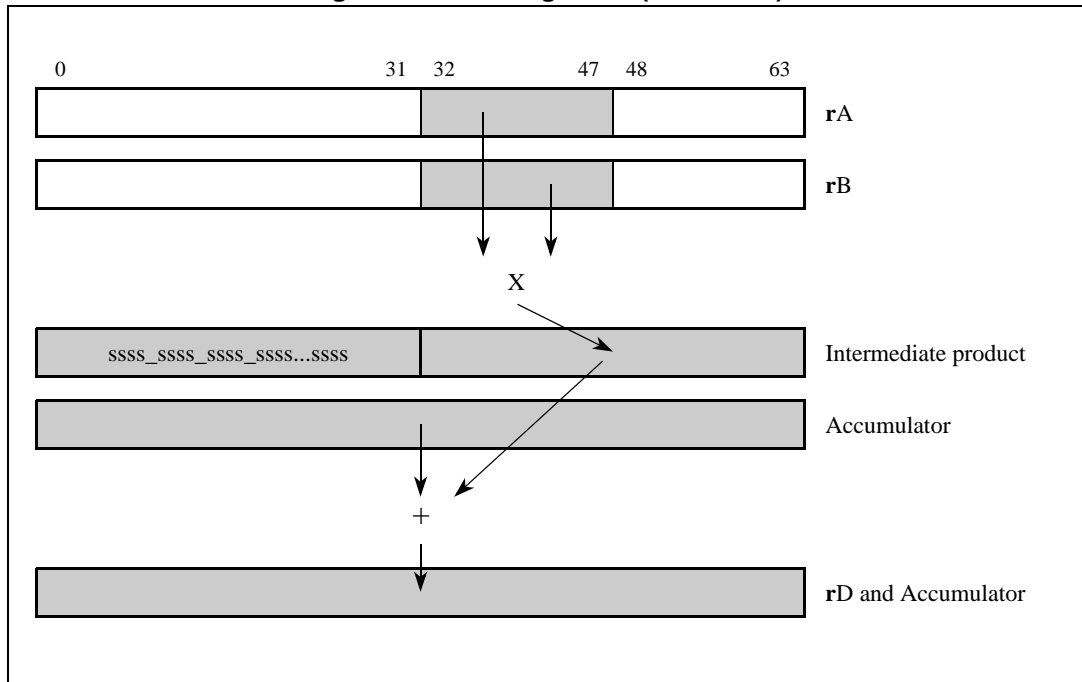
```
temp0:31 ← rA32:47 ×si rB32:47
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 + temp0:63

// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered half-word signed integer elements in **rA** and **rB** 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.

*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 136. *evmhegsmiaa* (even form)



**evmhegsmian** SPE APU User **evmhegsmian**  
**Vector multiply half words, even, guarded, signed, modulo, integer and accumulate negative**  
**evmhegsmian**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31							
0	0	0	1	0	0	rD										rA						rB						1	0	1	1	0	1	0	1	0	0	1

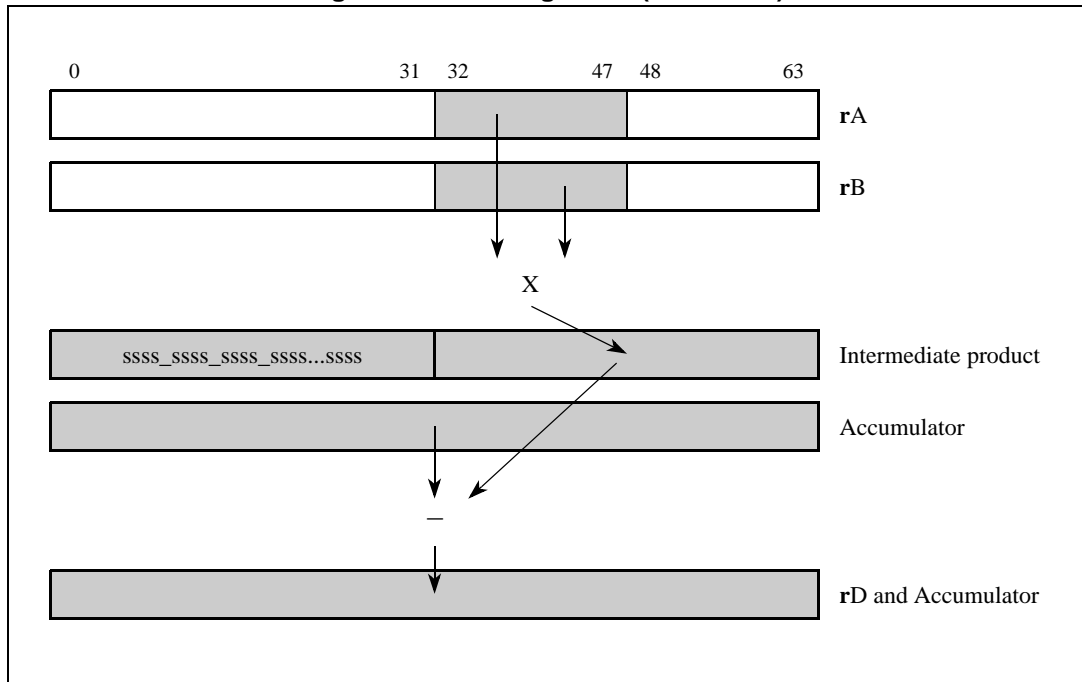
```
temp0:31 ← rA32:47 ×si rB32:47
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered half-word signed integer elements in **rA** and **rB** are multiplied. The intermediate product is sign-extended and subtracted from the contents of the 64-bit accumulator, and the result is placed into **rD** and into the accumulator.

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 137. evmhegsmian (even form)



evmhegumiaa

SPE APU	User
---------	------

evmhegumiaa

Vector multiply half words, even, guarded, unsigned, modulo, integer and accumulate

evmhegumiaa

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0					rD					rA					rB		1	0	1	0	0	1	0	1	0	0	0

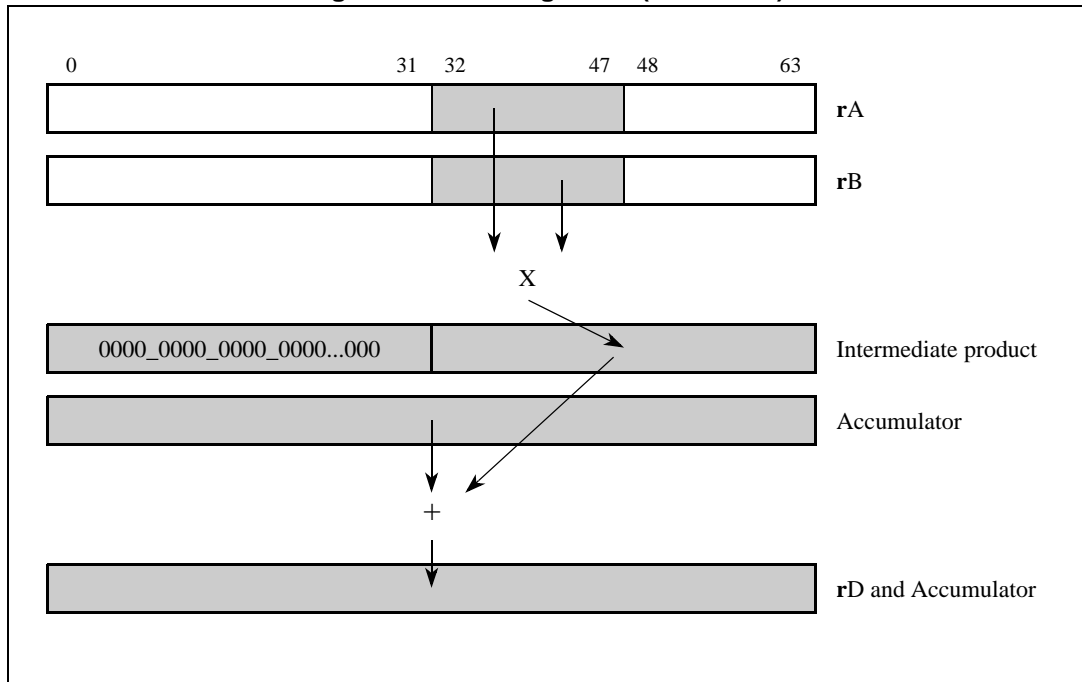
```
temp0:31 ← rA32:47 ×ui rB32:47
temp0:63 ← EXTZ(temp0:31)
rD0:63 ← ACC0:63 + temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered half-word unsigned integer elements in rA and rB 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.

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 138. *evmhегumiaа* (even form)



*evmhегumiaа*

SPE APU	User
---------	------

*evmhегumiaа*

Vector multiply half words, even, guarded, unsigned, modulo, integer and accumulate negative

*evmhегumiaа*                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0					rD					rA					rB					1	0	1	1	0	1	0	1	0	0	0


```
temp0:31 ← rA32:47 ×ui rB32:47
temp0:63 ← EXTZ(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low even-numbered unsigned integer elements in *rA* and *rB* are multiplied. The intermediate product is zero-extended and subtracted from the contents of the 64-bit accumulator. The result is placed into *rD* and into the accumulator.

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 139. evmhegumian (even form)



**evmhesmf** SPE APU User **evmhesmf**

**Vector multiply half words, even, signed, modulo, fractional (to accumulator)**

**evmhesmf** **rD,rA,rB** **(A = 0)**

**evmhesmfa** **rD,rA,rB** **(A = 1)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD										rA					rB					1	0	0	0	0	A	0	1	0	1	1

```
// high
rD0:31 ← (rA0:15 ×sf rB0:15)

// low
rD32:63 ← (rA32:47 ×sf rB32:47)

// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```


The corresponding even-numbered half-word signed fractional elements in rA and rB are multiplied then placed into the corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)



Figure 140. Even multiply of two signed modulo fractional elements (to accumulator) (evmhesmf)



**evmhesmfaaw** SPE APU User **evmhesmfaaw**

**Vector multiply half words, even, signed, modulo, fractional and accumulate into words**

**evmhesmfaaw** rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD										rA					rB					1	0	1	0	0	0	0	1	0	1	1

```

// high
temp0:31 ← (rA0:15 ×sf rB0:15)
rD0:31 ← ACC0:31 + temp0:31


// low
temp0:31 ← (rA32:47 ×sf rB32:47)
rD32:63 ← ACC32:63 + temp0:31

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding even-numbered half-word signed fractional elements in rA and rB 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.

Other registers altered: ACC

Figure 141. Even form of vector half-word multiply (evmhesmfaww)



evmhesmfaww                      SPE APU    User                      evmhesmfaww

**Vector multiply half words, even, signed, modulo, fractional and accumulate negative into words**

evmhesmfaww                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	0	0	0	1	0	1	1

```
// high
temp0:31 ← rA0:15 ×sf rB0:15
rD0:31 ← ACC0:31 - temp0:31


// low
temp0:31 ← rA32:47 ×sf rB32:47
rD32:63 ← ACC32:63 - temp0:31

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding even-numbered half-word signed fractional elements in rA and rB 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.

Other registers altered: ACC

Figure 142. Even form of vector half-word multiply (evmhesmfanw)



**evmhesmi** SPE APU User **evmhesmi**

**Vector multiply half words, even, signed, modulo, integer (to accumulator)**

**evmhesmi** **rD,rA,rB** **(A = 0)**

**evmhesmia** **rD,rA,rB** **(A = 1)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	0	A	0	1	0	0	1

```

// high
rD0:31 ← rA0:15 ×si rB0:15

// low
rD32:63 ← rA32:47 ×si rB32:47


// update accumulator
if A = 1, then ACC0:63 ← rD0:63
    
```

The corresponding even-numbered half-word signed integer elements in **rA** and **rB** are multiplied. The two 32-bit products are placed into the corresponding words of **rD**.

If **A = 1**, the result in **rD** is also placed into the accumulator.

Other registers altered: ACC (If **A = 1**)

Figure 143. Even form for vector multiply (to accumulator) (evmhesmi)



**evmhesmiaaw** SPE APU User **evmhesmiaaw**  
**Vector multiply half words, even, signed, modulo, integer and accumulate into words**  
**evmhesmiaaw** **rD,rA,rB**

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	0	0	0	1	0	0	rD	rA	rB	1	0	1	0	0	0	0	0	1	0	0	1											

```
// high
temp0:31 ← rA0:15 ×si rB0:15
rD0:31 ← ACC0:31 + temp0:31


// low
temp0:31 ← rA32:47 ×si rB32:47
rD32:63 ← ACC32:63 + temp0:31

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding even-numbered half-word signed integer elements in rA and rB 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.

Other registers altered: ACC

Figure 144. Even form of vector half-word multiply (evmhesmiaaw)



**evmhesmianw** SPE APU User **evmhesmianw**

**Vector multiply half words, even, signed, modulo, integer and accumulate negative into words**

**evmhesmianw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31															
0	0	0	1	0	0	rD										rA										rB										1	0	1	1	0	0	0	1	0	0	1

```

// high
temp0:31 ← rA0:15 ×si rB0:15
rD0:31 ← ACC0:31 - temp0:31


// low
temp1:31 ← rA32:47 ×si rB32:47
rD32:63 ← ACC32:63 - temp1:31

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding even-numbered half-word signed integer elements in **rA** and **rB** 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 **rD** words and into the accumulator.

Other registers altered: ACC

Figure 145. Even form of vector half-word multiply (evmhesmianw)



**evmhessf** SPE APU User **evmhessf**  
**Vector multiply half words, even, signed, saturate, fractional (to accumulator)**  
**evmhessf** **rD,rA,rB** **(A = 0)**  
**evmhessfa** **rD,rA,rB** **(A = 1)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD										rA					rB					1	0	0	0	0	A	0	0	0	1	1

```

// high
temp0:31 ← rA0:15 ×sf rB0:15
if (rA0:15 = 0x8000) & (rB0:15 = 0x8000) then
    rD0:31 ← 0x7FFF_FFFF //saturate
    movh ← 1
else
    rD0:31 ← temp0:31
    movh ← 0

// low
temp0:31 ← rA32:47 ×sf rB32:47
if (rA32:47 = 0x8000) & (rB32:47 = 0x8000) then
    rD32:63 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    rD32:63 ← temp0:31
    movl ← 0

// update accumulator
    
```



```

rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA32:47 ×sf rB32:47
if (rA32:47 = 0x8000) & (rB32:47 = 0x8000) then
    temp0:31 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    movl ← 0
temp0:63 ← EXTS(ACC32:63) + EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← movh
SPEFSCROV ← movl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh | movh
SPEFSCRSOV ← SPEFSCRSOV | ovl | movl
    
```

The corresponding even-numbered half-word signed fractional elements in **rA** and **rB** 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.

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 147. Even form of vector half-word multiply (evmhessfaaw)**





**evmhessfanw** SPE APU User **evmhessfanw**  
**Vector multiply half words, even, signed, saturate, fractional and accumulate negative into words**  
**evmhessfanw** **rD,rA,rB**




++The corresponding even-numbered half-word signed fractional elements in rA and rB 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.


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 148. Even form of vector half-word multiply (evmhessfanw)**



**evmhessiaaw** SPE APU User **evmhessiaaw**  
**Vector multiply half words, even, signed, saturate, integer and accumulate into words**  
**evmhessiaaw** **rD,rA,rB**



```
// high
temp0:31 ← rA0:15 ×si rB0:15
temp0:63 ← EXTS(ACC0:31) + EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
```

```

temp0:31 ← rA32:47 ×si rB32:47
temp0:63 ← EXTS(ACC32:63) + EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovl
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovl
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The corresponding even-numbered half-word signed integer elements in rA and rB 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.

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 149. Even form of vector half-word multiply (evmhessianw)**



**evmhessianw** SPE APU User **evmhessianw**

**Vector multiply half words, even, signed, saturate, integer and accumulate negative into words**

**evmhessianw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	0	0	0	0	0	0	0	0	0	1

```

// high
temp0:31 ← rA0:15 ×si rB0:15
temp0:63 ← EXTS(ACC0:31) - EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA32:47 ×si rB32:47
temp0:63 ← EXTS(ACC32:63) - EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63

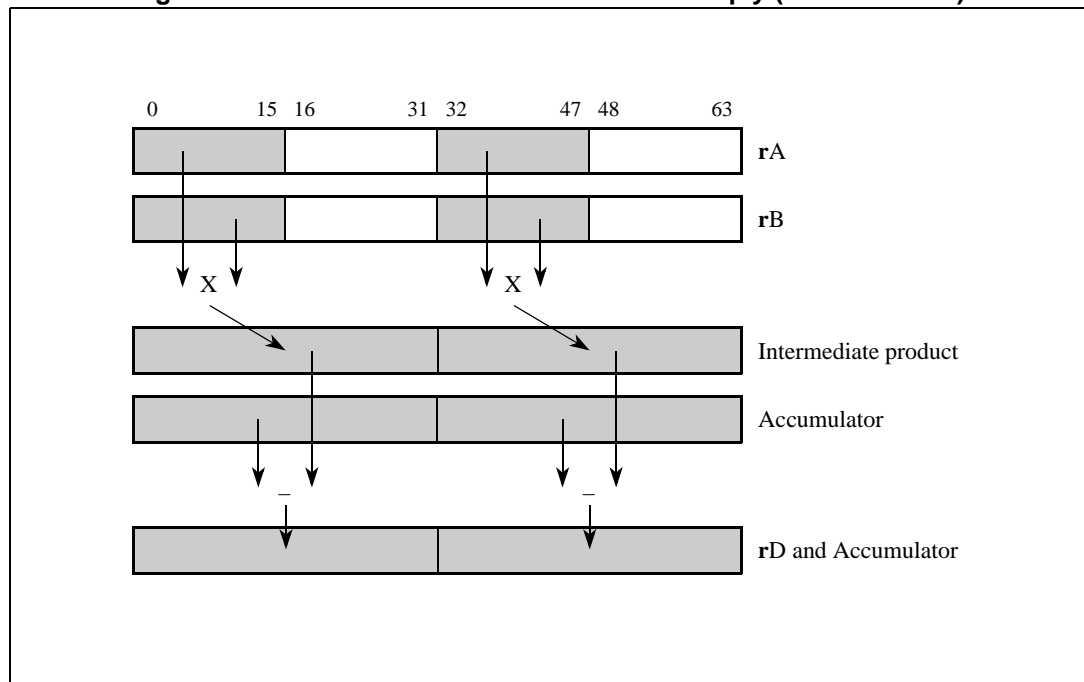
// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The corresponding even-numbered half-word signed integer elements in **rA** and **rB** 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.

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 150. Even form of vector half-word multiply (evmhessianw)**




**evmheumi** SPE APU User **evmheumi**

**Vector multiply half words, even, unsigned, modulo, integer (to accumulator)**

**evmheumi** **rD,rA,rB** **(A = 0)**

**evmheumia** **rD,rA,rB** **(A = 1)**



```
// high
rD0:31 ← rA0:15 ×ui rB0:15


// low
rD32:63 ← rA32:47 ×ui rB32:47

// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding even-numbered half-word unsigned integer elements in rA and rB are multiplied. The two 32-bit products are placed into the corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.


**Figure 151. Vector multiply half words, even, unsigned, modulo, integer (to accumulator) (evmheumi)**



**evmheumiaaw** SPE APU User **evmheumiaaw**

**Vector multiply half words, even, unsigned, modulo, integer and accumulate into words**

**evmheumiaaw** **rD,rA,rB**



```
// high
temp0:31 ← rA0:15 ×ui rB0:15
```

```

rD0:31 ← ACC0:31 + temp0:31


// low
temp0:31 ← rA32:47 ×ui rB32:47
rD32:63 ← ACC32:63 + temp0:31

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding even-numbered half-word unsigned integer elements in rA and rB 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.

Other registers altered: ACC

**Figure 152. Even form of vector half-word multiply (evmheumiaaw)**



**evmheumianw**  
**evmheumianw**

SPE APU	User
---------	------

**Vector multiply half words, even, unsigned, modulo, integer and accumulate negative into words**

**evmheumianw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																											
										<b>rD</b>					<b>rA</b>					<b>rB</b>												
																				1 0 1 1 0 0 0 1 0 0 0												

```

// high
temp0:31 ← rA0:15 ×ui rB0:15
rD0:31 ← ACC0:31 - temp0:31
    
```


```
// low
temp0:31 ← rA32:47 ×ui rB32:47
rD32:63 ← ACC32:63 - temp0:31

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding even-numbered half-word unsigned integer elements in rA and rB are multiplied. Each intermediate product is subtracted from the contents of the corresponding accumulator words. The differences are placed into the corresponding rD and accumulator words.

Other registers altered: ACC

Figure 153. Even form of vector half-word multiply (evmheumianw)



**evmheusiaaw** SPE APU | User **evmheusiaaw**

**Vector multiply half words, even, unsigned, saturate, integer and accumulate into words**

**evmheusiaaw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31																					
0	0	0	1	0	0	rD										rA										rB										1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0

```
// high
temp0:31 ← rA0:15 ×ui rB0:15
temp0:63 ← EXTZ(ACC0:31) + EXTZ(temp0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)
```

```

//low
temp0:31 ← rA32:47 ×ui rB32:47
temp0:63 ← EXTZ(ACC32:63) + EXTZ(temp0:31)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

For each word element in the accumulator, corresponding even-numbered half-word unsigned integer elements in *rA* and *rB* 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.

If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

**Figure 154. Even form of vector half-word multiply (evmheusiaaw)**




**evmheusianw**

SPE APU	User
---------	------

**evmheusianw**

**Vector multiply half words, even, unsigned, saturate, integer and accumulate negative into words**

**evmheusianw                    rD,rA,rB**



```

// high
temp0:31 ← rA0:15 ×ui rB0:15
temp0:63 ← EXTZ(ACC0:31) - EXTZ(temp0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0x0000_0000, 0x0000_0000, temp32:63)

//low
temp0:31 ← rA32:47 ×ui rB32:47
temp0:63 ← EXTZ(ACC32:63) - EXTZ(temp0:31)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0x0000_0000, 0x0000_0000, temp32:63)

// update accumulator
ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

For each word element in the accumulator, corresponding even-numbered half-word unsigned integer elements in **rA** and **rB** 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.


If there is an underflow from the subtraction, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC





Figure 155. Even form of vector half-word multiply (evmheusianw)



evmhogsmfaa

SPE APU	User
---------	------

evmhogsmfaa

Vector multiply half words, odd, guarded, signed, modulo, fractional and accumulate

evmhogsmfaa

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD										rA					rB					1	0	1	0	0	1	0	1	1	1	1


$$\begin{aligned} \text{temp}_{0:31} &\leftarrow rA_{48:63} \times_{sf} rB_{48:63} \\ \text{temp}_{0:63} &\leftarrow \text{EXTS}(\text{temp}_{0:31}) \\ rD_{0:63} &\leftarrow \text{ACC}_{0:63} + \text{temp}_{0:63} \end{aligned}$$

// update accumulator  
 $\text{ACC}_{0:63} \leftarrow rD_{0:63}$

The corresponding low odd-numbered half-word signed fractional elements in rA and rB 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.

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 156. evmhogsmfaa (odd form)



evmhogsmfan

SPE APU	User
---------	------

evmhogsmfan

Vector multiply half words, odd, guarded, signed, modulo, fractional and accumulate negative

evmhogsmfan                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																	1	0	1	1	0	1	0	1	1	1	1


```
temp0:31 ← rA48:63 ×sf rB48:63
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low odd-numbered half-word signed fractional elements in rA and rB 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.

*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 157. *evmhogsmfan* (odd form)



**evmhogsmiaa**

SPE APU	User
---------	------

**evmhogsmiaa**

**Vector multiply half words, odd, guarded, signed, modulo, integer, and accumulate**

**evmhogsmiaa**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0																1	0	1	0	0	1	0	1	1	0	1
						rD						rA						rB													

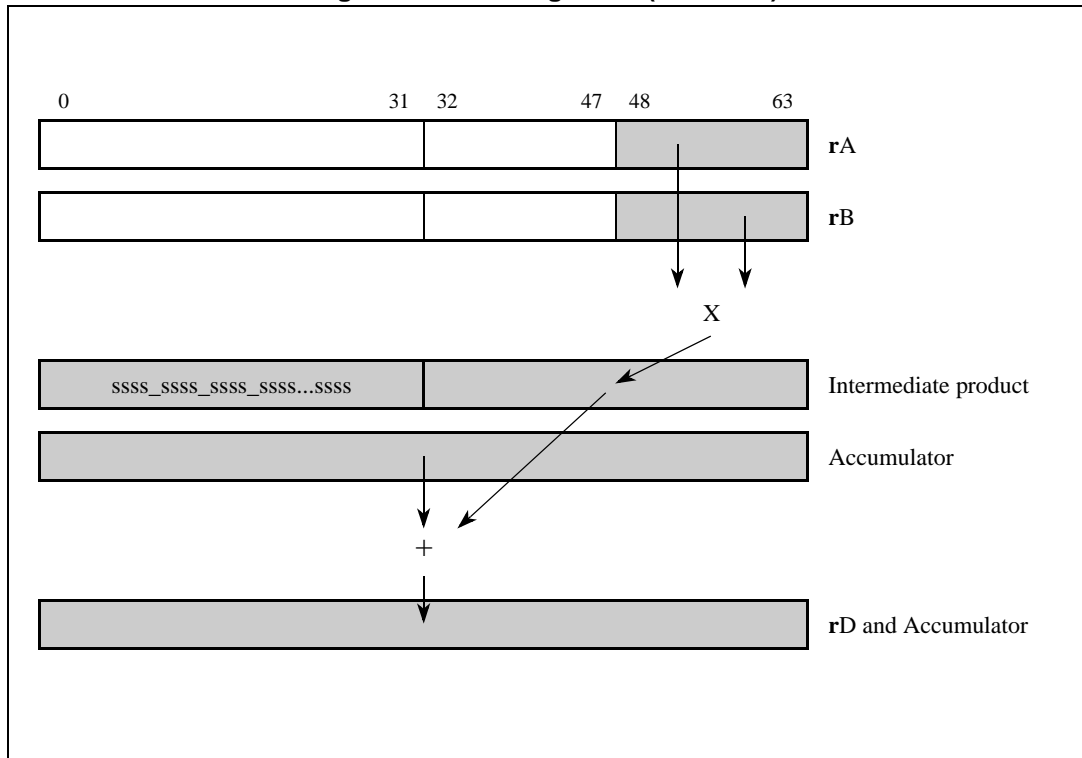
```
temp0:31 ← rA48:63 ×si rB48:63
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 + temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low odd-numbered half-word signed integer elements in **rA** and **rB** 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.

*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 158. evmhogsmiaa (odd form)



evmhogsmian SPE APU User evmhogsmian

**Vector multiply half words, odd, guarded, signed, modulo, integer and accumulate negative**

evmhogsmian rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	0	1	0	1	1	0	1


```
temp0:31 ← rA48:63 ×si rB48:63
temp0:63 ← EXTS(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low odd-numbered half-word signed integer elements in rA and rB 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.

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 159. evmhogsmian (odd form)



**evmhogumiaa** SPE APU | User **evmhogumiaa**  
**Vector multiply half words, odd, guarded, unsigned, modulo, integer and accumulate**  
**evmhogumiaa** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD										rA					rB					1	0	1	0	0	1	0	1	1	0	0

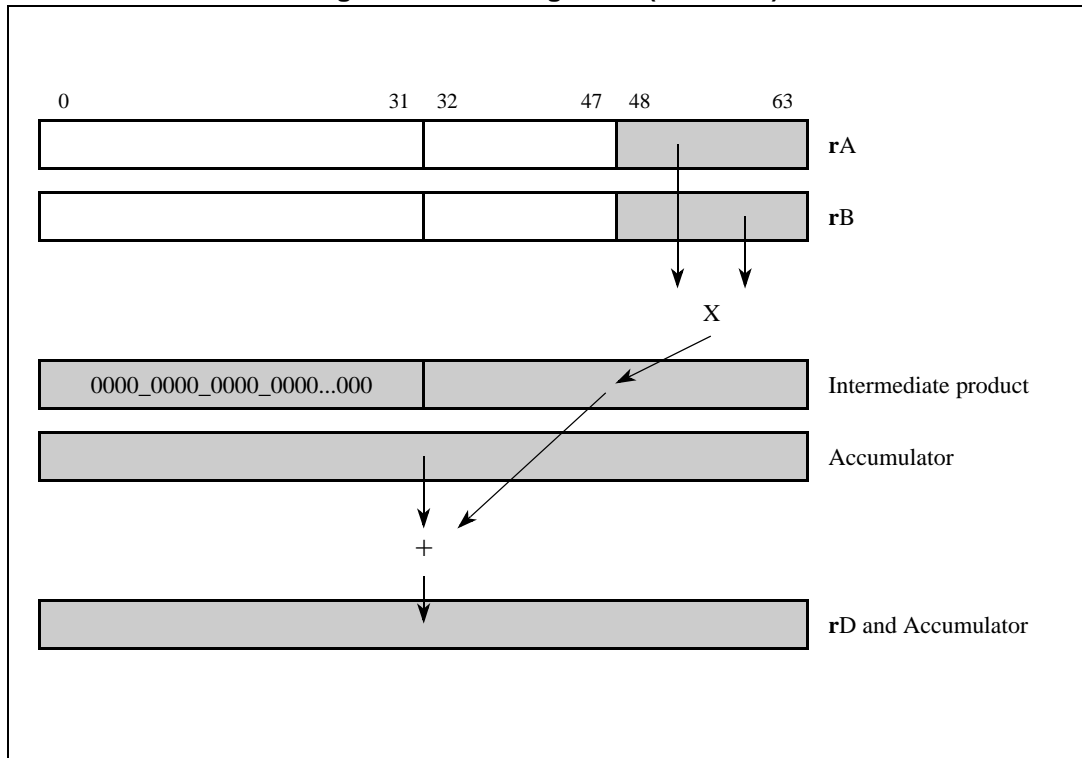
$$\begin{aligned} \text{temp}_{0:31} &\leftarrow rA_{48:63} \times_{ui} rB_{48:63} \\ \text{temp}_{0:63} &\leftarrow \text{EXTZ}(\text{temp}_{0:31}) \\ rD_{0:63} &\leftarrow \text{ACC}_{0:63} + \text{temp}_{0:63} \end{aligned}$$

// update accumulator  
 $\text{ACC}_{0:63} \leftarrow rD_{0:63}$

The corresponding low odd-numbered half-word unsigned integer elements in rA and rB 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.

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 160. evmhogumiaa (odd form)



evmhogumian

SPE APU	User
---------	------

evmhogumian

Vector multiply half words, odd, guarded, unsigned, modulo, integer and accumulate negative

evmhogumian                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																											
						rD						rA						rB						1 0 1 1 0 1 0 1 1 0 0								


```
temp0:31 ← rA48:63 ×ui rB48:63
temp0:63 ← EXTZ(temp0:31)
rD0:63 ← ACC0:63 - temp0:63
```

```
// update accumulator
ACC0:63 ← rD0:63
```

The corresponding low odd-numbered half-word unsigned integer elements in rA and rB 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.

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 161. evmhogumian (odd form)



**evmhosmf** SPE APU User **evmhosmf**

**Vector multiply half words, odd, signed, modulo, fractional (to accumulator)**

**evmhosmf** **rD,rA,rB** **(A = 0)**

**evmhosmfa** **rD,rA,rB** **(A = 1)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	0	A	0	1	1	1	1

```
// high
rD0:31 ← (rA16:31 ×sf rB16:31)

// low
rD32:63 ← (rA48:63 ×sf rB48:63)


// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding odd-numbered, half-word signed fractional elements in rA and rB are multiplied. Each product is placed into the corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

Figure 162. Vector multiply half words, odd, signed, modulo, fractional (to accumulator) (evmhosmf)



**evmhosmfaaw** SPE APU User **evmhosmfaaw**

**Vector multiply half words, odd, signed, modulo, fractional and accumulate into words**

**evmhosmfaaw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	0	0	0	0	1	1	1	1

```

// high
temp0:31 ← rA16:31 ×sf rB16:31
rD0:31 ← ACC0:31 + temp0:31

// low
temp0:31 ← rA48:63 ×sf rB48:63
rD32:63 ← ACC32:63 + temp0:31


// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed fractional elements in rA and rB 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 rD words and into the accumulator

Other registers altered: ACC



Figure 163. Odd form of vector half-word multiply (evmhosmfaaw)



evmhosmfanw

SPE APU	User
---------	------

evmhosmfanw

Vector multiply half words, odd, signed, modulo, fractional and accumulate negative into words

evmhosmfanw

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																	1	0	1	1	0	0	0	1	1	1	1

```

// high
temp0:31 ← rA16:31 ×sf rB16:31
rD0:31 ← ACC0:31 - temp0:31


// low
temp0:31 ← rA48:63 ×sf rB48:63
rD32:63 ← ACC32:63 - temp0:31

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed fractional elements in rA and rB 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 rD words and into the accumulator.

Other registers altered: ACC

Figure 164. Odd form of vector half-word multiply (evmhosmfanw)



**evmhosmi**

SPE APU	User
---------	------

**evmhosmi**

Vector multiply half words, odd, signed, modulo, integer (to accumulator)

<b>evmhosmi</b>	rD,rA,rB	(A = 0)
<b>evmhosmia</b>	rD,rA,rB	(A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0					rD				rA				rB			1	0	0	0	0	A	0	1	1	0	1

```
// high
rD0:31 ← rA16:31 ×si rB16:31

// low
rD32:63 ← rA48:63 ×si rB48:63


// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding odd-numbered half-word signed integer elements in rA and rB are multiplied. The two 32-bit products are placed into the corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

Figure 165. Vector multiply half words, odd, signed, modulo, integer (to accumulator) (evmhosmi)



evmhosmiaaw

SPE APU	User
---------	------

evmhosmiaaw

Vector multiply half words, odd, signed, modulo, integer and accumulate into words

evmhosmiaaw

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	0	0	0	0	1	1	0	1

```
// high
temp0:31 ← rA16:31 ×si rB16:31
rD0:31 ← ACC0:31 + temp0:31


// low
temp0:31 ← rA48:63 ×si rB48:63
rD32:63 ← ACC32:63 + temp0:31

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed integer elements in rA and rB 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 rD words and into the accumulator.

Other registers altered: ACC

Figure 166. Odd form of vector half-word multiply (evmhosmiaaw)



**evmhosmianw**

SPE APU | User

**evmhosmianw**

**Vector multiply half words, odd, signed, modulo, integer and accumulate negative into words**

**evmhosmianw**

**rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																	1	0	1	1	0	0	0	1	1	0	1
										rD					rA					rB												

```

// high
temp0:31 ← rA16:31 ×si rB16:31
rD0:31 ← ACC0:31 - temp0:31

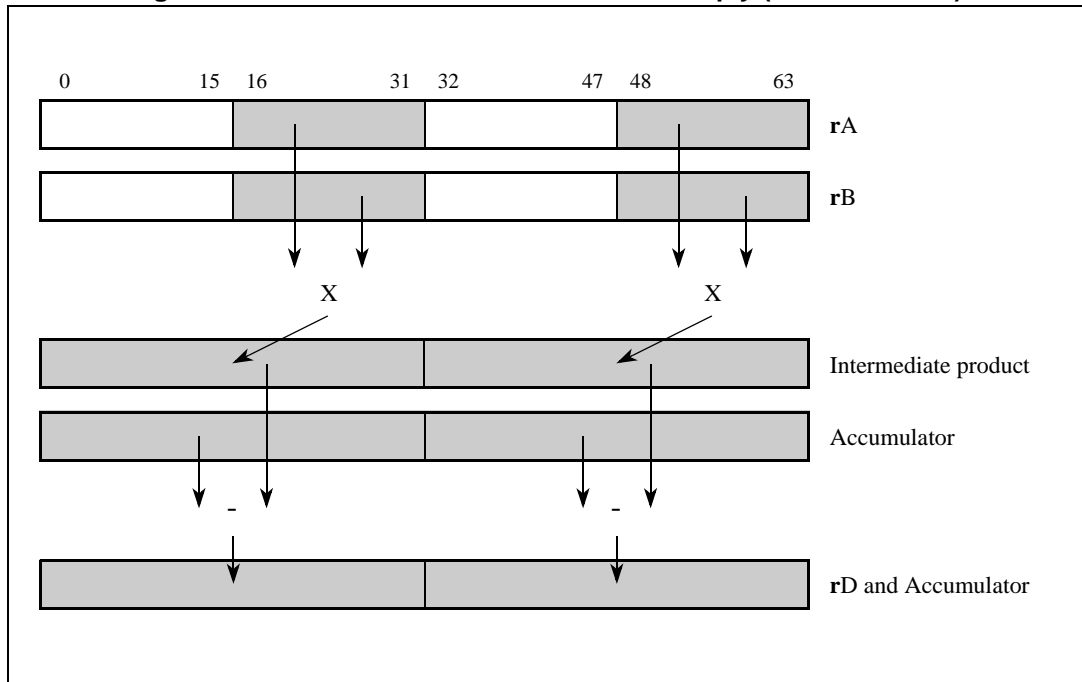
// low
temp0:31 ← rA48:63 ×si rB48:63
rD32:63 ← ACC32:63 - temp0:31

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding odd-numbered half-word signed integer elements in `rA` and `rB` 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 `rD` words and into the accumulator.

Other registers altered: ACC

Figure 167. Odd form of vector half-word multiply (evmhosmianw)



**evmhossf** SPE APU User **evmhossf**  
**Vector multiply half words, odd, signed, saturate, fractional (to accumulator)**  
**evmhossf** **rD,rA,rB** **(A = 0)**  
**evmhossfa** **rD,rA,rB** **(A = 1)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	0	A	0	0	1	1	1

```

// high
temp0:31 ← rA16:31 ×sf rB16:31
if (rA16:31 = 0x8000) & (rB16:31 = 0x8000) then
    rD0:31 ← 0x7FFF_FFFF //saturate
    movh ← 1
else
    rD0:31 ← temp0:31
    movh ← 0

// low
temp0:31 ← rA48:63 ×sf rB48:63
if (rA48:63 = 0x8000) & (rB48:63 = 0x8000) then
    rD32:63 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    rD32:63 ← temp0:31
    movl ← 0

// update accumulator
    
```

```

if A = 1 then ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← movh
SPEFSCROV ← movl
SPEFSCRSOVH ← SPEFSCRSOVH | movh
SPEFSCRSOV ← SPEFSCRSOV | movl


```

The corresponding odd-numbered half-word signed fractional elements in rA and rB are multiplied. The 32 bits of each product are placed into the corresponding words of rD. 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 A = 1, the result in rD is also placed into the accumulator.

Other registers altered: SPEFSCR  
ACC (If A = 1)

**Figure 168. Vector multiply half words, odd, signed, saturate, fractional (to accumulator) (evmhossf)**



evmhossfaaw

SPE APU	User
---------	------

evmhossfaaw

**Vector multiply half words, odd, signed, saturate, fractional and accumulate into words**

evmhossfaaw rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0											rD																

```

// high
temp0:31 ← rA16:31 ×sf rB16:31

```



```

if (rA16:31 = 0x8000) & (rB16:31 = 0x8000) then
    temp0:31 ← 0x7FFF_FFFF //saturate
    movh ← 1
else
    movh ← 0
temp0:63 ← EXTS(ACC0:31) + EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA48:63 ×sf rB48:63
if (rA48:63 = 0x8000) & (rB48:63 = 0x8000) then
    temp0:31 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    movl ← 0
temp0:63 ← EXTS(ACC32:63) + EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← movh
SPEFSCROV ← movl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh | movh
SPEFSCRSOV ← SPEFSCRSOV | ovl | movl


```

The corresponding odd-numbered half-word signed fractional elements in **rA** and **rB** 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.

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 169. Odd form of vector half-word multiply (evmhossfaaw)



evmhossfanw

SPE APU	User
---------	------

evmhossfanw

Vector multiply half words, odd, signed, saturate, fractional and accumulate negative into words

evmhossfanw                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																											

```

// high
temp0:31 ← rA16:31 ×sf rB16:31
if (rA16:31 = 0x8000) & (rB16:31 = 0x8000) then
    temp0:31 ← 0x7FFF_FFFF //saturate
    movh ← 1
else
    movh ← 0
temp0:63 ← EXTS(ACC0:31) - EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA48:63 ×sf rB48:63
if (rA48:63 = 0x8000) & (rB48:63 = 0x8000) then
    temp0:31 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    movl ← 0
temp0:63 ← EXTS(ACC32:63) - EXTS(temp0:31)
    
```



```

ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← movh
SPEFSCROV ← movl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh | movh
SPEFSCRSOV ← SPEFSCRSOV | ovl | movl
    
```

The corresponding odd-numbered half-word signed fractional elements in **rA** and **rB** 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.

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 170. odd Form of Vector Half-Word Multiply (evmhossfanw)**




evmhossiaaw

SPE APU	User
---------	------

evmhossiaaw

Vector multiply half words, odd, signed, saturate, integer and accumulate into words

evmhossiaaw                    rD,rA,rB



```

// high
temp0:31 ← rA16:31 ×si rB16:31
temp0:63 ← EXTS(ACC0:31) + EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA48:63 ×si rB48:63
temp0:63 ← EXTS(ACC32:63) + EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The corresponding odd-numbered half-word signed integer elements in rA and rB 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.

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 171. Odd form of vector half-word multiply (evmhossiaaw)



**evmhossianw**

SPE APU	User
---------	------

**evmhossianw**

**Vector multiply half words, odd, signed, saturate, integer and accumulate negative into words**

**evmhossianw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	0	0	0	0	0	1	0	1

```

// high
temp0:31 ← rA16:31 ×si rB16:31
temp0:63 ← EXTS(ACC0:31) - EXTS(temp0:31)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:31 ← rA48:63 ×si rB48:63
temp0:63 ← EXTS(ACC32:63) - EXTS(temp0:31)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
    
```


$$\begin{aligned} \text{SPEFSCR}_{\text{SOVH}} &\leftarrow \text{SPEFSCR}_{\text{SOVH}} \mid \text{ovh} \\ \text{SPEFSCR}_{\text{SOV}} &\leftarrow \text{SPEFSCR}_{\text{SOV}} \mid \text{ovl} \end{aligned}$$

The corresponding odd-numbered half-word signed integer elements in rA and rB 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.

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 172. Odd form of vector half-word multiply (evmhossianw)



evmhoumi evmhoumi

SPE APU | User

Vector multiply half words, odd, unsigned, modulo, integer (to accumulator)

evmhoumi rD,rA,rB (A = 0)  
 evmhoumia rD,rA,rB (A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	0	A	0	1	1	0	0

```
// high
rD0:31 ← rA16:31 ×ui rB16:31

// low
rD32:63 ← rA48:63 ×ui rB48:63


// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding odd-numbered half-word unsigned integer elements in **rA** and **rB** are multiplied. The two 32-bit products are placed into the corresponding words of **rD**.

If **A = 1**, the result in **rD** is also placed into the accumulator.

Other registers altered: **ACC** (If **A = 1**)

**Figure 173. Vector multiply half words, odd, unsigned, modulo, integer (to accumulator) (evmhomi)**



**evmhomiaaw** SPE APU User **evmhomiaaw**

**Vector multiply half words, odd, unsigned, modulo, integer and accumulate into words**

**evmhomiaaw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	0	0	0	0	1	1	0	0

```
// high
temp0:31 ← rA16:31 ×ui rB16:31
rD0:31 ← ACC0:31 + temp0:31

// low
temp0:31 ← rA48:63 ×ui rB48:63
rD32:63 ← ACC32:63 + temp0:31


// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding odd-numbered half-word unsigned integer elements in **rA** and **rB** 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.

Other registers altered: ACC

**Figure 174. Odd form of vector half-word multiply (evmhoumiaaw)**



**evmhoumianw**

SPE APU	User
---------	------

**evmhoumianw**

**Vector multiply half words, odd, unsigned, modulo, integer and accumulate negative into words**

**evmhoumianw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																	1	0	1	1	0	0	0	1	1	0	0

```

// high
temp0:31 ← rA0:15 ×ui rB0:15
rD0:31 ← ACC0:31 - temp0:31
/
/ low
temp0:31 ← rA32:47 ×ui rB32:47
rD32:63 ← ACC32:63 - temp0:31


// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding odd-numbered half-word unsigned integer elements in rA and rB are multiplied. Each intermediate product is

subtracted from the contents of the corresponding accumulator word. The results are placed into the corresponding rD and accumulator words.

Other registers altered: ACC

**Figure 175. Odd form of vector half-word multiply (evmhoumianw)**



**evmhouisiaaw** SPE APU User **evmhouisiaaw**

**Vector multiply half words, odd, unsigned, saturate, integer and accumulate into words**

**evmhouisiaaw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31													
0	0	0	1	0	0	rD									rA									rB									1	0	1	0	0	0	0	0	0	1	0	0

```

// high
temp0:31 ← rA16:31 ×ui rB16:31
temp0:63 ← EXTZ(ACC0:31) + EXTZ(temp0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)

//low
temp0:31 ← rA48:63 ×ui rB48:63
temp0:63 ← EXTZ(ACC32:63) + EXTZ(temp0:31)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)
    
```

```
// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
```

For each word element in the accumulator, corresponding odd-numbered half-word unsigned integer elements in **rA** and **rB** 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.

If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

**Figure 176. Odd form of vector half-word multiply (evmhouisiaaw)**






evmhousianw

SPE APU	User
---------	------

evmhousianw

Vector multiply half words, odd, unsigned, saturate, integer and accumulate negative into words

evmhousianw                    rD,rA,rB



```

// high
temp0:31 ← rA16:31 ×ui rB16:31
temp0:63 ← EXTZ(ACC0:31) - EXTZ(temp0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)

//low
temp0:31 ← rA48:63 ×ui rB48:63
temp0:63 ← EXTZ(ACC32:63) - EXTZ(temp0:31)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

For each word element in the accumulator, corresponding odd-numbered half-word unsigned integer elements in rA and rB 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.

If subtraction causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

Figure 177. Odd form of vector half-word multiply (evmhousesianw)



evmra

SPE APU	User
---------	------

evmra

Initialize accumulator

evmra

rD,rA

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	1	0	0	1	1	0	0	0	1	0	0


$$ACC_{0:63} \leftarrow rA_{0:63}$$

$$rD_{0:63} \leftarrow rA_{0:63}$$

The contents of rA are written into the accumulator and copied into rD. This is the method for initializing the accumulator.

Other registers altered: ACC

Figure 178. Initialize accumulator (evmra)




**evmwhsmf**

SPE APU	User
---------	------

**evmwhsmf**

**Vector multiply word high signed, modulo, fractional (to accumulator)**

**evmwhsmf**  $rD, rA, rB$  (A = 0)  
**evmwhsmfa**  $rD, rA, rB$  (A = 1)



```

// high
temp0:63 ← rA0:31 ×sf rB0:31
rD0:31 ← temp0:31

// low
temp0:63 ← rA32:63 ×sf rB32:63
rD32:63 ← temp0:31


// update accumulator
if A = 1 then ACC0:63 ← rD0:63
    
```

The corresponding word signed fractional elements in rA and rB are multiplied and bits 0–31 of the two products are placed into the two corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (if A =1)

**Figure 179. Vector multiply word high signed, modulo, fractional (to accumulator) (evmwhsmf)**




**evmwhsmi**

SPE APU	User
---------	------

**evmwhsmi**

**Vector multiply word high signed, modulo, integer (to accumulator)**

**evmwhsmi**  $rD, rA, rB$  (A = 0)  
**evmwhsmia**  $rD, rA, rB$  (A = 1)



```
// high
temp0:63 ← rA0:31 ×si rB0:31
rD0:31 ← temp0:31

// low
temp0:63 ← rA32:63 ×si rB32:63
rD32:63 ← temp0:31


// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding word signed integer elements in **rA** and **rB** are multiplied. Bits 0–31 of the two 64-bit products are placed into the two corresponding words of **rD**.

If **A = 1**, The result in **rD** is also placed into the accumulator.

Other registers altered: ACC (If **A = 1**)

**Figure 180. Vector multiply word high signed, modulo, integer (to accumulator) (evmwhsm)**



**evmwhssf**

SPE APU	User
---------	------

**evmwhssf**

**Vector multiply word high signed, saturate, fractional (to accumulator)**

**evmwhssf** **rD,rA,rB** (A = 0)  
**evmwhssfA** **rD,rA,rB** (A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	0	0	1	1	1

```

// high
temp0:63 ← rA0:31 ×sf rB0:31
if (rA0:31 = 0x8000_0000) & (rB0:31 = 0x8000_0000) then
    rD0:31 ← 0x7FFF_FFFF //saturate
    movh ← 1
else
    rD0:31 ← temp0:31
    movh ← 0

// low
temp0:63 ← rA32:63 ×sf rB32:63
if (rA32:63 = 0x8000_0000) & (rB32:63 = 0x8000_0000) then
    rD32:63 ← 0x7FFF_FFFF //saturate
    movl ← 1
else
    rD32:63 ← temp0:31
    movl ← 0


// update accumulator
if A = 1 then ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← movh
SPEFSCROV ← movl
SPEFSCRSOVH ← SPEFSCRSOVH | movh
SPEFSCRSOV ← SPEFSCRSOV | movl
    
```

The corresponding word signed fractional elements in **rA** and **rB** are multiplied. Bits 0–31 of each product are placed into the corresponding words of **rD**. 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 181. Vector multiply word high signed, saturate, fractional (to accumulator) (evmwhssf)



**evmwhumi**

SPE APU	User
---------	------

**evmwhumi**

Vector multiply word high unsigned, modulo, integer (to accumulator)

**evmwhumi** **rD,rA,rB** (A = 0)  
**evmwhumia** **rD,rA,rB** (A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	0	1	1	0	0

```
// high
temp0:63 ← rA0:31 ×ui rB0:31
rD0:31 ← temp0:31

// low
temp0:63 ← rA32:63 ×ui rB32:63
rD32:63 ← temp0:31


// update accumulator
if A = 1, ACC0:63 ← rD0:63
```

The corresponding word unsigned integer elements in rA and rB are multiplied. Bits 0–31 of the two products are placed into the two corresponding words of rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

Figure 182. Vector multiply word high unsigned, modulo, integer (to accumulator) (evmwhumi)



**evmwlsmiaaw** SPE APU | User **evmwlsmiaaw**  
**Vector multiply word low signed, modulo, integer and accumulate in words**  
**evmwlsmiaaw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	1	0	0	1	0	0	1

```
// high
temp0:63 ← rA0:31 ×si rB0:31
rD0:31 ← ACC0:31 + temp32:63

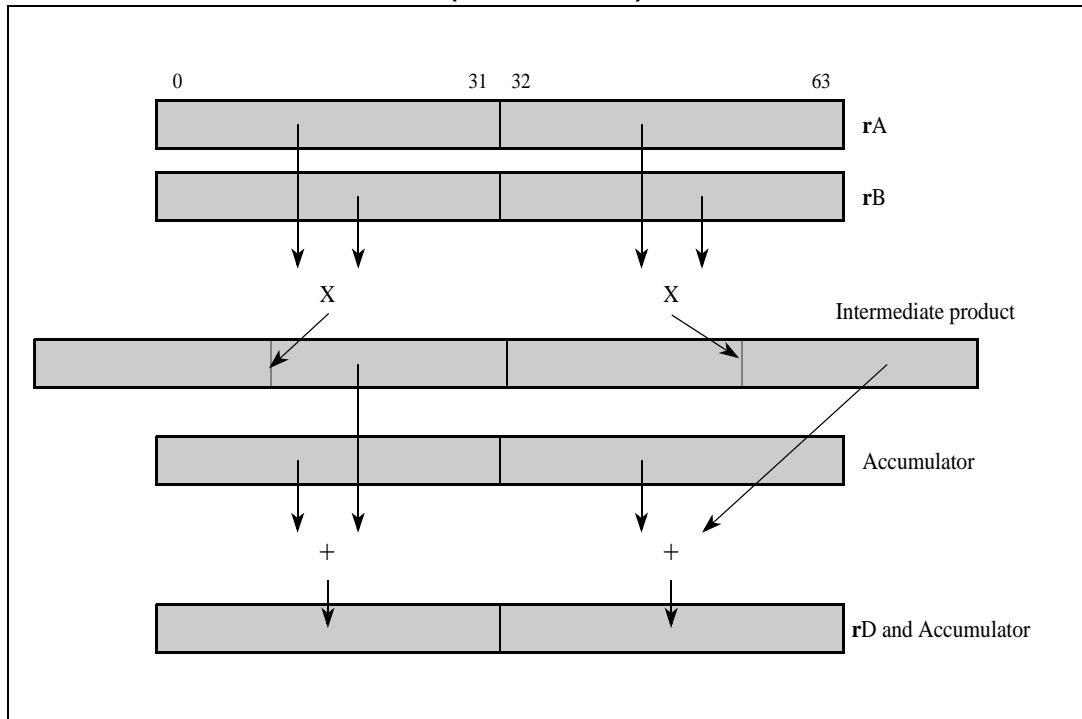
// low
temp0:63 ← rA32:63 ×si rB32:63
rD32:63 ← ACC32:63 + temp32:63

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding word signed integer elements in rA and rB 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 rD and the accumulator.

Other registers altered: ACC

Figure 183. Vector multiply word low signed, modulo, integer & accumulate in words (evmwlsmiaaw)



evmwlsnianw

SPE APU	User
---------	------

evmwlsnianw

Vector multiply word low signed, modulo, integer and accumulate negative in words

evmwlsnianw                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0																										

```

// high
temp0:63 ← rA0:31 ×si rB0:31
rD0:31 ← ACC0:31 - temp32:63

// low
temp0:63 ← rA32:63 ×si rB32:63
rD32:63 ← ACC32:63 - temp32:63


// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding word elements in rA and rB 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 rD and the accumulator.

Other registers altered: ACC



Figure 184. Vector multiply word low signed, modulo, integer and accumulate negative in words (evmwlsianw)



evmwlsiaaw

SPE APU	User
---------	------

evmwlsiaaw

Vector multiply word low signed, saturate, integer and accumulate in words

evmwlsiaaw                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																	1	0	1	0	1	0	0	0	0	0	1

```

// high
temp0:63 ← rA0:31 ×si rB0:31
temp0:63 ← EXTS(ACC0:31) + EXTS(temp32:63)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:63 ← rA32:63 ×si rB32:63
temp0:63 ← EXTS(ACC32:63) + EXTS(temp32:63)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← ovh
    
```


$SPEFSCR_{OV} \leftarrow ov_l$   
 $SPEFSCR_{SOVH} \leftarrow SPEFSCR_{SOVH} | ov_h$   
 $SPEFSCR_{SOV} \leftarrow SPEFSCR_{SOV} | ov_l$

The corresponding word signed integer elements in **rA** and **rB** are multiplied producing a 64-bit product. The 32 lsbs of each product is added to the corresponding word in the ACC, saturating if overflow or underflow occurs; the result is placed in **rD** and the accumulator.

If there is overflow or underflow from the addition, overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

**Figure 185. Vector multiply word low signed, saturate, integer & accumulate in words (evmlssiaaw)**



**evmlssianw** SPE APU User **evmlssianw**

**Vector multiply word low signed, saturate, integer and accumulate negative in words**

**evmlssianw** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	1	0	0	0	0	0	0	0	1

```

// high
temp0:63 ← rA0:31 ×si rB0:31
temp0:63 ← EXTS(ACC0:31) - EXTS(temp32:63)
ovh ← (temp31 ⊕ temp32)
rD0:31 ← SATURATE(ovh, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// low
temp0:63 ← rA32:63 ×si rB32:63
    
```

```

temp0:63 ← EXTS(ACC32:63) - EXTS(temp32:63)
ovl ← (temp31 ⊕ temp32)
rD32:63 ← SATURATE(ovl, temp31, 0x8000_0000, 0x7FFF_FFFF, temp32:63)

// update accumulator
ACC0:63 ← rD0:63


// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

The corresponding word signed integer elements in rA and rB 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.


If addition causes overflow or underflow, overflow and summary overflow SPEFSCR bits are recorded.

Other registers altered: SPEFSCR ACC

**Figure 186. Vector multiply word low signed, saturate, integer & accumulate negative in words (evmwlsianw)**



<b>evmwlumi</b>	SPE APU	User	<b>evmwlumi</b>
<b>Vector multiply word low unsigned, modulo, integer</b>			
<b>evmwlumi</b>	<b>rD,rA,rB</b>		(A = 0)
<b>evmwlumia</b>	<b>rD,rA,rB</b>		(A = 1)



```

// high
temp0:63 ← rA0:31 ×ui rB0:31
rD0:31 ← temp32:63

// low
temp0:63 ← rA32:63 ×ui rB32:63
rD32:63 ← temp32:63

// update accumulator
If A = 1 then ACC0:63 ← rD0:63
    
```

The corresponding word unsigned integer elements in **rA** and **rB** are multiplied. The least significant 32 bits of each product are placed into the two corresponding words of **rD**.


Note: The least significant 32 bits of the product are independent of whether the word elements in **rA** and **rB** are treated as signed or unsigned 32-bit integers.

If A = 1, the result in **rD** 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 187. Vector multiply word low unsigned, modulo, integer (evmwlumi)**



evmwlumiaaw

SPE APU	User
---------	------

evmwlumiaaw

Vector multiply word low unsigned, modulo, integer and accumulate in words

evmwlumiaaw                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						rA						rB						1	0	1	0	1	0	0	1	0	0	0	0

```

// high
temp0:63 ← rA0:31 ×ui rB0:31
rD0:31 ← ACC0:31 + temp32:63


// low
temp0:63 ← rA32:63 ×ui rB32:63
rD32:63 ← ACC32:63 + temp32:63

// update accumulator
ACC0:63 ← rD0:63
    
```

For each word element in the accumulator, the corresponding word unsigned integer elements in rA and rB 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 rD and the accumulator.

Other registers altered: ACC

**Figure 188. Vector multiply word low unsigned, modulo, integer & accumulate in words (evmwlumiaaw)**



evmwlumianw

SPE APU	User
---------	------

evmwlumianw

Vector multiply word low unsigned, modulo, integer and accumulate negative in words

evmwlumianw                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						1	0	1	1	1	0	0	1	0	0	0

```
// high
temp0:63 ← rA0:31 ×ui rB0:31
rD0:31 ← ACC0:31 - temp32:63

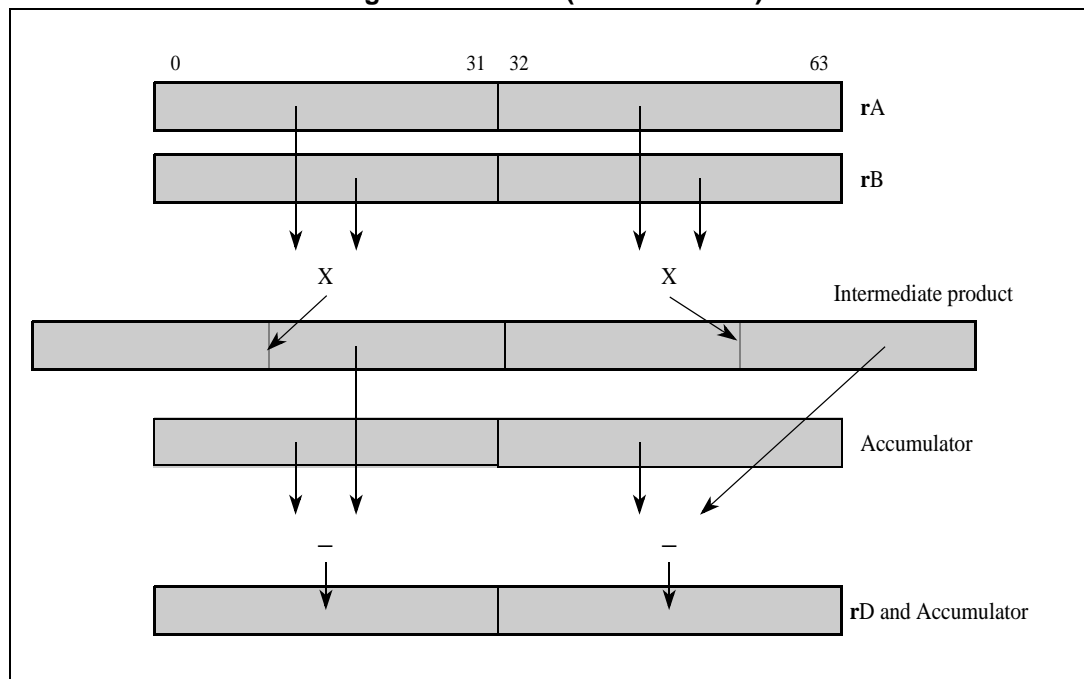
// low
temp0:63 ← rA32:63 ×ui rB32:63
rD32:63 ← ACC32:63 - temp32:63

// update accumulator
ACC0:63 ← rD0:63
```

For each word element in the accumulator, the corresponding word unsigned integer elements in rA and rB 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 rD and the accumulator.

Other registers altered: ACC

Figure 189. Vector multiply word low unsigned, modulo, integer & accumulate negative in words (evmwlumianw)



evmwlusiaaw

SPE APU	User
---------	------

evmwlusiaaw

Vector multiply word low unsigned, saturate, integer and accumulate in words

evmwlusiaaw

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD						rA						rB						1	0	1	0	1	0	0	0	0	0	0	0	0

```

// high
temp0:63 ← rA0:31 ×ui rB0:31
temp0:63 ← EXTZ(ACC0:31) + EXTZ(temp32:63)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)

//low
temp0:63 ← rA32:63 ×ui rB32:63
temp0:63 ← EXTZ(ACC32:63) + EXTZ(temp32:63)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0xFFFF_FFFF, 0xFFFF_FFFF, temp32:63)


// update accumulator
ACC0:63 ← rD0:63
// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

For each word element in the ACC, corresponding word unsigned integer elements in rA and rB 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.

If the addition causes overflow, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

Figure 190. Vector multiply word low unsigned, saturate, integer & accumulate in words (evmwlusiaaw)



**evmwlusianw**

SPE APU	User
---------	------

**evmwlusianw**

**Vector multiply word low unsigned, saturate, integer and accumulate negative in words**

**evmwlusianw**                      **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	1	0	0	0	0	0	0	0	0	0	0

```
// high
temp0:63 ← rA0:31 ×ui rB0:31
temp0:63 ← EXTZ(ACC0:31) - EXTZ(temp32:63)
ovh ← temp31
rD0:31 ← SATURATE(ovh, 0, 0x0000_0000, 0x0000_0000, temp32:63)

//low
temp0:63 ← rA32:63 ×ui rB32:63
temp0:63 ← EXTZ(ACC32:63) - EXTZ(temp32:63)
ovl ← temp31
rD32:63 ← SATURATE(ovl, 0, 0x0000_0000, 0x0000_0000, temp32:63)
// update accumulator
ACC0:63 ← rD0:63
// update SPEFSCR
SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
```




For each ACC word element, corresponding word elements in rA and rB 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 rD and the ACC.

If there is an underflow from the subtraction, the overflow and summary overflow bits are recorded in the SPEFSCR.

Other registers altered: SPEFSCR ACC

**Figure 191. Vector multiply word low unsigned, saturate, integer & accumulate negative in words (evmwlsuanw)**



<b>evmwsmf</b>	SPE APU   User	<b>evmwsmf</b>
<b>Vector multiply word signed, modulo, fractional (to accumulator)</b>		
<b>evmwsmf</b>	<b>rD,rA,rB</b>	<b>(A = 0)</b>
<b>evmwsmfa</b>	<b>rD,rA,rB</b>	<b>(A = 1)</b>

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	1	1	0	1	1

$$rD_{0:63} \leftarrow rA_{32:63} \times_{sf} rB_{32:63}$$


```
// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The corresponding low word signed fractional elements in rA and rB are multiplied. The product is placed into rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

Figure 192. Vector multiply word signed, modulo, fractional (to accumulator) (evmwsmf)



evmwsmfaa SPE APU User evmwsmfaa  
**Vector multiply word signed, modulo, fractional and accumulate**  
 evmwsmfaa rD,rA,rB

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	0	0	0	1	0	0	rD	rA	rB	1	0	1	0	1	0	1	1	0	1	1	0	1	1									

$$\text{temp}_{0:63} \leftarrow rA_{32:63} \times_{sf} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} + \text{temp}_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The corresponding low word signed fractional elements in rA and rB are multiplied. The intermediate product is added to the contents of the 64-bit accumulator and the result is placed in rD and the accumulator.

Other registers altered: ACC

Figure 193. Vector multiply word signed, modulo, fractional & accumulate (evmwsmlfaa)



evmwsmlfaa

SPE APU	User
---------	------

evmwsmlfaa

Vector multiply word signed, modulo, fractional and accumulate negative

evmwsmlfaa rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0					rD				rA						rB			1	0	1	1	1	0	1	1	0	1	1

$$\text{temp}_{0:63} \leftarrow rA_{32:63} \times_{sf} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} - \text{temp}_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The corresponding low word signed fractional elements in rA and rB are multiplied. The intermediate product is subtracted from the contents of the accumulator and the result is placed in rD and the accumulator.

Other registers altered: ACC

Figure 194. Vector multiply word signed, modulo, fractional & accumulate negative (evmwsmfan)



<b>evmwsmi</b>	SPE APU	User	<b>evmwsmi</b>
<b>Vector multiply word signed, modulo, integer (to accumulator)</b>			
<b>evmwsmi</b>	<b>rD,rA,rB</b>		(A = 0)
<b>evmwsmia</b>	<b>rD,rA,rB</b>		(A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	1	1	0	0	1

$$rD_{0:63} \leftarrow rA_{32:63} \times_{si} rB_{32:63}$$


```
// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The low word signed integer elements in rA and rB are multiplied. The product is placed into rD.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

Figure 195. Vector multiply word signed, modulo, integer (to accumulator) (evmwsmi)



evmwsmiaa SPE APU User evmwsmiaa

**Vector multiply word signed, modulo, integer and accumulate**

evmwsmiaa                      rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	1	0	1	1	0	0	1

$$\text{temp}_{0:63} \leftarrow rA_{32:63} \times_{si} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} + \text{temp}_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The low word signed integer elements in rA and rB are multiplied. The intermediate product is added to the contents of the 64-bit accumulator and the result is placed into rD and the accumulator.

Other registers altered: ACC

Figure 196. Vector multiply word signed, modulo, integer & accumulate (evmwsmiaa)



evmwsnian SPE APU User evmwsnian

Vector multiply word signed, modulo, integer and accumulate negative

evmwsnian rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	1	0	1	1	0	0	1

$$temp_{0:63} \leftarrow rA_{32:63} \times_{si} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} - temp_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The low word signed integer elements in rA and rB are multiplied. The intermediate product is subtracted from the contents of the 64-bit accumulator and the result is placed into rD and the accumulator.

Other registers altered: ACC

Figure 197. Vector multiply word signed, modulo, integer & accumulate negative (evmwsmian)



<b>evmwssf</b>	SPE APU   User	<b>evmwssf</b>
<b>Vector multiply word signed, saturate, fractional (to accumulator)</b>		
<b>evmwssf</b>	<b>rD,rA,rB</b>	(A = 0)
<b>evmwssfA</b>	<b>rD,rA,rB</b>	(A = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	1	0	0	1	1

```

temp0:63 ← rA32:63 ×sf rB32:63
if (rA32:63 = 0x8000_0000) & (rB32:63 = 0x8000_0000) then
    rD0:63 ← 0x7FFF_FFFF_FFFF_FFFF //saturate
    mov ← 1
else
    rD0:63 ← temp0:63
    mov ← 0

// update accumulator
if A = 1 then ACC0:63 ← rD0:63

// update SPEFSCR
SPEFSCROVH ← 0
SPEFSCROV ← mov
SPEFSCRSOV ← SPEFSCRSOV | mov
    
```


The low word signed fractional elements in rA and rB are multiplied. The 64 bit product is placed into rD. If both inputs are -1.0, the result saturates to the largest positive 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 A = 1, the result in rD is also placed into the accumulator.

Other registers altered: SPEFSCR ACC (If A = 1)

**Figure 198. Vector multiply word signed, saturate, fractional (to accumulator) (evmwssf)**



evmwssfaa SPE APU User evmwssfaa

**Vector multiply word signed, saturate, fractional and accumulate**

evmwssfaa rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31							
0	0	0	1	0	0	rD										rA						rB						1	0	1	0	1	0	1	0	0	1	1

```

temp0:63 ← rA32:63 ×sf rB32:63
if (rA32:63 = 0x8000_0000) & (rB32:63 = 0x8000_0000) then
    temp0:63 ← 0x7FFF_FFFF_FFFF_FFFF //saturate
    mov ← 1
else
    mov ← 0
temp0:64 ← EXTS(ACC0:63) + EXTS(temp0:63)
ov ← (temp0 ⊕ temp1)
rD0:63 ← temp1:64 )
// update accumulator
ACC0:63 ← rD0:63
// update SPEFSCR
SPEFSCROVH ← 0
SPEFSCROV ← mov
SPEFSCRSOV ← SPEFSCRSOV | ov | mov
    
```




The low word signed fractional elements in rA and rB 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 rD and the ACC.

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 199. Vector multiply word signed, saturate, fractional, & accumulate (evmwssfaa)**



**evmwssfan** SPE APU User **evmwssfan**

**Vector multiply word signed, saturate, fractional and accumulate negative**

**evmwssfan** **rD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31															
0	0	0	1	0	0	rD										rA										rB										1	0	1	1	1	0	1	0	0	1	1

```

temp0:63 ← rA32:63 ×sf rB32:63
if (rA32:63 = 0x8000_0000) & (rB32:63 = 0x8000_0000) then
    temp0:63 ← 0x7FFF_FFFF_FFFF_FFFF //saturate
    mov ← 1
else
    mov ← 0
temp0:64 ← EXTS(ACC0:63) - EXTS(temp0:63)
ov ← (temp0 ⊕ temp1)
rD0:63 ← temp1:64 )
// update accumulator
ACC0:63 ← rD0:63
// update SPEFSCR
SPEFSCROVH ← 0
    
```

```
SPEFSCROV ← mov
SPEFSCRSOV ← SPEFSCRSOV | ov | mov
```


The low word signed fractional elements in **rA** and **rB** 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 **rD** and the ACC.

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 200. Vector multiply word signed, saturate, fractional & accumulate negative (evmwssfan)**



**evmwumi** **evmwumi**

SPE APU	User
---------	------

**Vector multiply word unsigned, modulo, integer (to accumulator)**

```
evmwumi          rD,rA,rB          (A = 0)
evmwumia        rD,rA,rB          (A = 1)
```

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rD					rA					rB					1	0	0	0	1	A	1	1	0	0	0	0

$$rD_{0:63} \leftarrow rA_{32:63} \times_{ui} rB_{32:63}$$


```
// update accumulator
if A = 1 then ACC0:63 ← rD0:63
```

The low word unsigned integer elements in **rA** and **rB** are multiplied to form a 64-bit product that is placed into **rD**.

If A = 1, the result in rD is also placed into the accumulator.

Other registers altered: ACC (If A = 1)

**Figure 201. Vector multiply word unsigned, modulo, integer (to accumulator) (evmwumi)**



evmwumiaa

SPE APU	User
---------	------

evmwumiaa

Vector multiply word unsigned, modulo, integer and accumulate

evmwumiaa

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	0	1	0	1	1	0	0	0

$$\text{temp}_{0:63} \leftarrow rA_{32:63} \times_{ui} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} + \text{temp}_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The low word unsigned integer elements in rA and rB 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.

Other registers altered: ACC

Figure 202. Vector multiply word unsigned, modulo, integer & accumulate (evmwumiaa)



evmwumian SPE APU User evmwumian  
**Vector multiply word unsigned, modulo, integer and accumulate negative**  
 evmwumian rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					1	0	1	1	1	0	1	1	0	0	0

$$\text{temp}_{0:63} \leftarrow rA_{32:63} \times_{ui} rB_{32:63}$$


$$rD_{0:63} \leftarrow ACC_{0:63} - \text{temp}_{0:63}$$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

The low word unsigned integer elements in **rA** and **rB** 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**.

Other registers altered: ACC

Figure 203. Vector multiply word unsigned, modulo, integer & accumulate negative (evmwumian)



evnand

SPE APU	User
---------	------

evnand

Vector NAND

evnand

rD,rA,rB

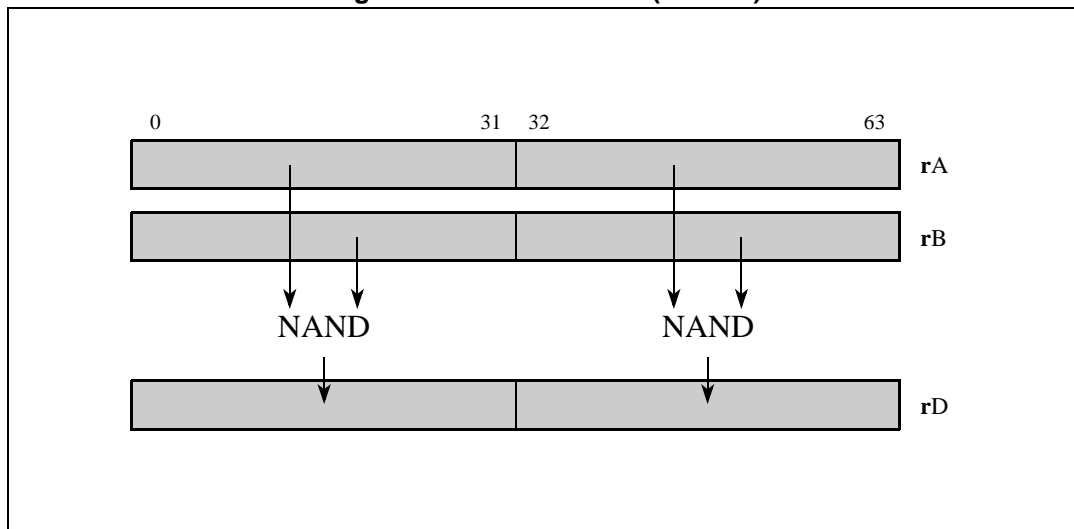
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					0	1	0	0	0	0	1	1	1	1	0

$$rD_{0:31} \leftarrow \neg(rA_{0:31} \& rB_{0:31}) \quad // \text{Bitwise NAND}$$

$$rD_{32:63} \leftarrow \neg(rA_{32:63} \& rB_{32:63}) \quad // \text{Bitwise NAND}$$

Corresponding word elements of **rA** and **rB** are bitwise Nanded. The result is placed in the corresponding element of **rD**.

Figure 204. Vector NAND (evnand)



**evneg** SPE APU User **evneg**  
**Vector negate**  
**evneg** rD,rA


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	1

$$rD_{0:31} \leftarrow \text{NEG}(rA_{0:31})$$

$$rD_{32:63} \leftarrow \text{NEG}(rA_{32:63})$$

The negative of each element of rA is placed in rD. The negative of 0x8000\_0000 (most negative number) returns 0x8000\_0000. No overflow is detected.

Figure 205. Vector negate (evneg)



**evnor**

SPE APU	User
---------	------

**evnor**

**Vector NOR**

**evnor**  $rD, rA, rB$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	0	1	1	0	0	0


$$rD_{0:31} \leftarrow \neg(rA_{0:31} \mid rB_{0:31}) \quad // \text{Bitwise NOR}$$

$$rD_{32:63} \leftarrow \neg(rA_{32:63} \mid rB_{32:63}) \quad // \text{Bitwise NOR}$$

Each element of **rA** and **rB** is bitwise NORed. The result is placed in the corresponding element of **rD**.

*Note:* Use **evnand** or **evnor** for **evnot**.

**Figure 206. Vector NOR (evnor)**



Simplified mnemonic: **evnot rD,rA** performs a complement register

**evnot rD,rA** equivalent to **evnor rD,rA,rA**

**evor**

SPE APU	User
---------	------

**evor**

**Vector OR**

**evor**  $rD, rA, rB$

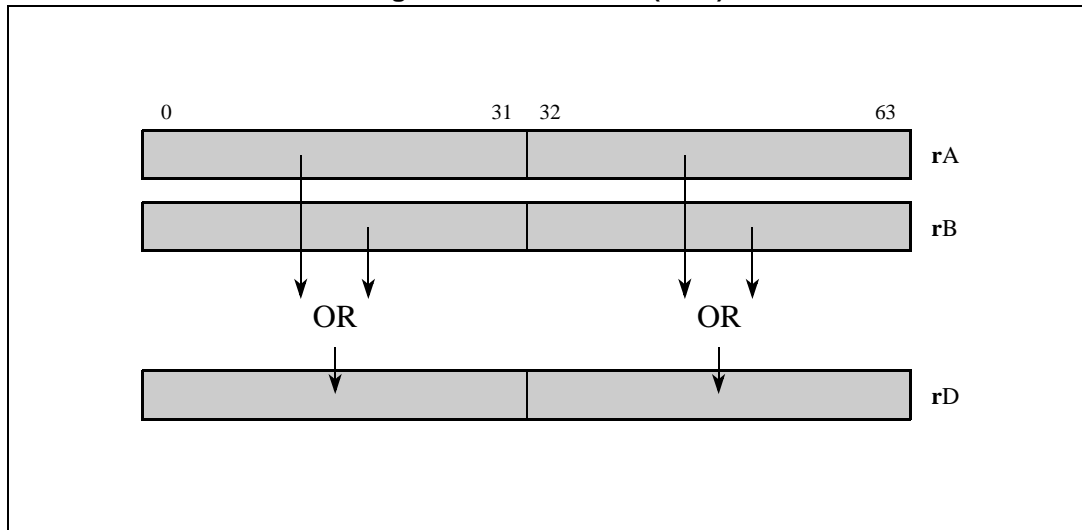
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	0	1	0	1	1	1

$$rD_{0:31} \leftarrow rA_{0:31} \mid rB_{0:31} \quad // \text{Bitwise OR}$$

$$rD_{32:63} \leftarrow rA_{32:63} \mid rB_{32:63} \quad // \text{Bitwise OR}$$

Each element of **rA** and **rB** is bitwise ORed. The result is placed in the corresponding element of **rD**.

Figure 207. 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 APU	User
---------	------

**evorc**

**Vector OR with complement**

**evorc**

rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0																											
						rD						rA						rB														

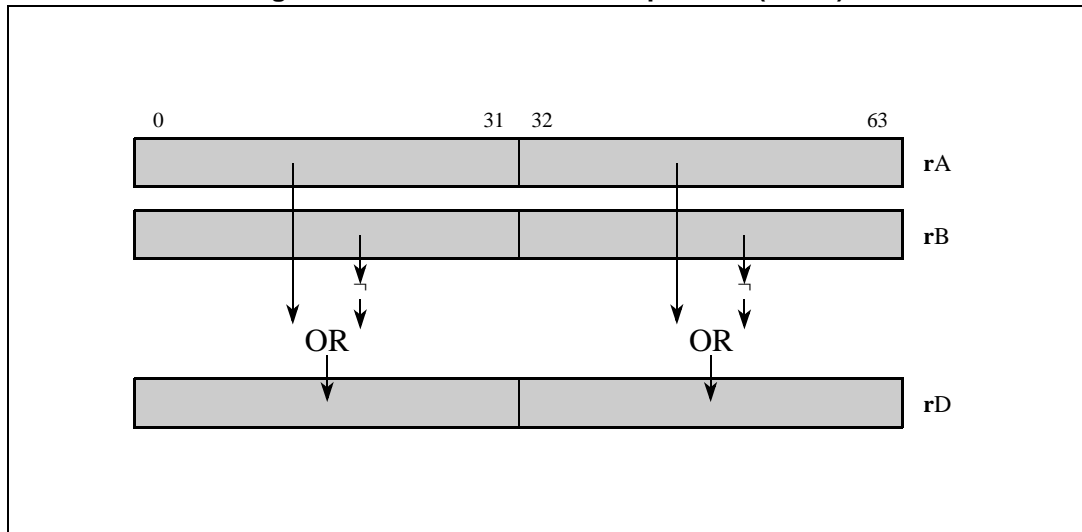
$$rD_{0:31} \leftarrow rA_{0:31} | (\neg rB_{0:31}) \quad // \text{ Bitwise ORC}$$

$$rD_{32:63} \leftarrow rA_{32:63} | (\neg rB_{32:63}) \quad // \text{ Bitwise ORC}$$

Each element of rA is bitwise ORed with the complement of rB. The result is placed in the corresponding element of rD.



Figure 208. Vector OR with complement (evorc)



evrlw

SPE APU	User
---------	------

evrlw

Vector rotate left word

evrlw


rD,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	1	0	1	0	0	0

$$\begin{aligned}
 nh &\leftarrow rB_{27:31} \\
 nl &\leftarrow rB_{59:63} \\
 rD_{0:31} &\leftarrow ROTL(rA_{0:31}, nh) \\
 rD_{32:63} &\leftarrow ROTL(rA_{32:63}, nl)
 \end{aligned}$$

Each of the high and low elements of rA is rotated left by an amount specified in rB. The result is placed into rD. Rotate values for each element of rA are found in bit positions rB[27–31] and rB[59–63].

Figure 209. Vector rotate left word (evrlw)



**evrlwi** SPE APU User **evrlwi**  
**Vector rotate left word immediate**  
**evrlwi** **rD,rA,UIMM**

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
	0	0	0	1	0	0	<b>rD</b>					<b>rA</b>					<b>UIMM</b>					0	1	0	0	0	1	0	1	0	1	0


$$n \leftarrow \text{UIMM}$$

$$rD_{0:31} \leftarrow \text{ROTL}(rA_{0:31}, n)$$

$$rD_{32:63} \leftarrow \text{ROTL}(rA_{32:63}, n)$$

Both the high and low elements of rA are rotated left by an amount specified by a 5-bit immediate value.

Figure 210. Vector rotate left word immediate (evrlwi)



**evrndw**

SPE APU	User
---------	------

**evrndw**

**Vector round word**

**evrndw**

**rD,rA**


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD					rA					0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	0	0

```

rD0:31 ← (rA0:31+0x00008000) & 0xFFFF0000 // Modulo sum
rD32:63 ← (rA32:63+0x00008000) & 0xFFFF0000 // Modulo sum
    
```

The 32-bit elements of rA are rounded into 16 bits. The result is placed into rD. The resulting 16 bits are placed in the most significant 16 bits of each element of rD, zeroing out the low order 16 bits of each element.

**Figure 211. Vector round word (evrndw)**



**evsel**

SPE APU	User
---------	------

**evsel**

**Vector select**

**evsel**

**rD,rA,rB,crS**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	0	0	1	0	0	rD					rA					rB					0	1	0	0	1	1	1	1	crS				


```

ch ← CRcrS*4
cl ← CRcrS*4+1
if (ch = 1) then rD0:31 ← rA0:31
else rD0:31 ← rB0:31
if (cl = 1) then rD32:63 ← rA32:63
else rD32:63 ← rB32:63
    
```

If the most significant bit in the crS field of CR is set, the high-order element of rA is placed in the high-order element of rD; otherwise, the high-order element of rB is placed into the high-order element of rD. If the next most significant bit in the crS field of CR is set, the low-

order element of **rA** is placed in the low-order element of **rD**, otherwise, the low-order element of **rB** is placed into the low-order element of **rD**. This is shown in [Figure 212](#).

**Figure 212. Vector select (evsel)**



**evslw**

SPE APU	User
---------	------

**evslw**

**Vector shift left word**

**evslw**

**rD,rA,rB**

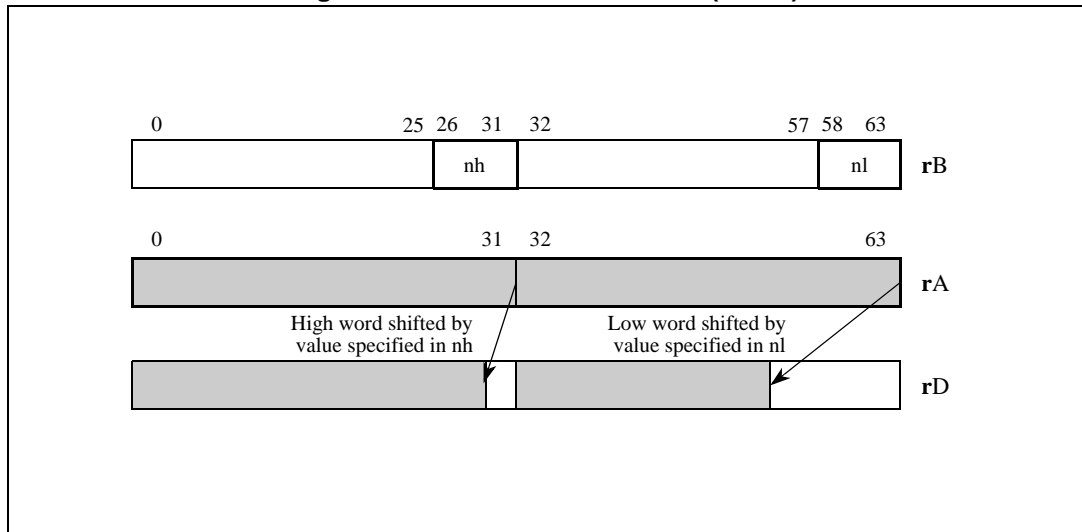
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	<b>rD</b>						<b>rA</b>						<b>rB</b>						0	1	0	0	0	1	0	0	1	0	0

$$\begin{aligned}
 nh &\leftarrow rB_{26:31} \\
 nl &\leftarrow rB_{58:63} \\
 rD_{0:31} &\leftarrow SL(rA_{0:31}, nh) \\
 rD_{32:63} &\leftarrow SL(rA_{32:63}, nl)
 \end{aligned}$$

Each of the high and low elements of **rA** are shifted left by an amount specified in **rB**. The result is placed into **rD**. 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 213. Vector shift left word (evslw)



**evslwi** SPE APU User **evslwi**

**Vector shift left word immediate**

**evslwi** **rD,rA,UIMM**



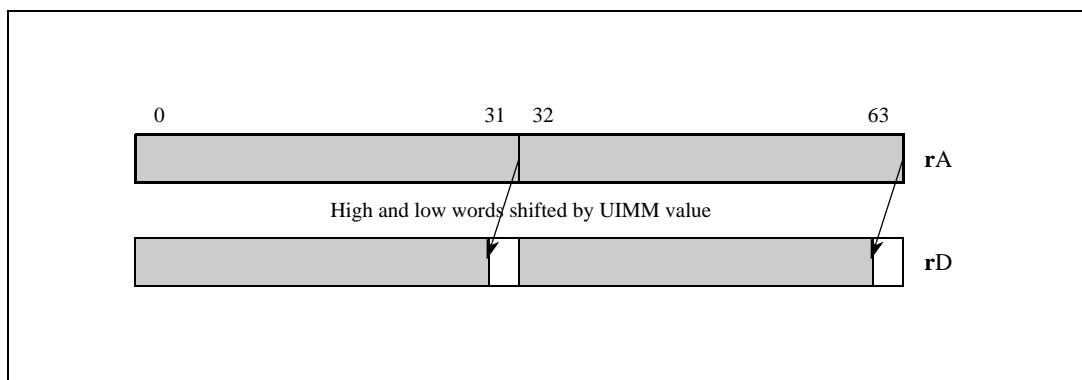
$$n \leftarrow \text{UIMM}$$

$$rD_{0:31} \leftarrow \text{SL}(rA_{0:31}, n)$$

$$rD_{32:63} \leftarrow \text{SL}(rA_{32:63}, n)$$

Both high and low elements of rA are shifted left by the 5-bit UIMM value and the results are placed in rD.

Figure 214. Vector shift left word immediate (evslwi)




**evsplatfi**

SPE APU	User
---------	------

**evsplatfi**

**Vector splat fractional immediate**

**evsplatfi**                      rD,SIMM

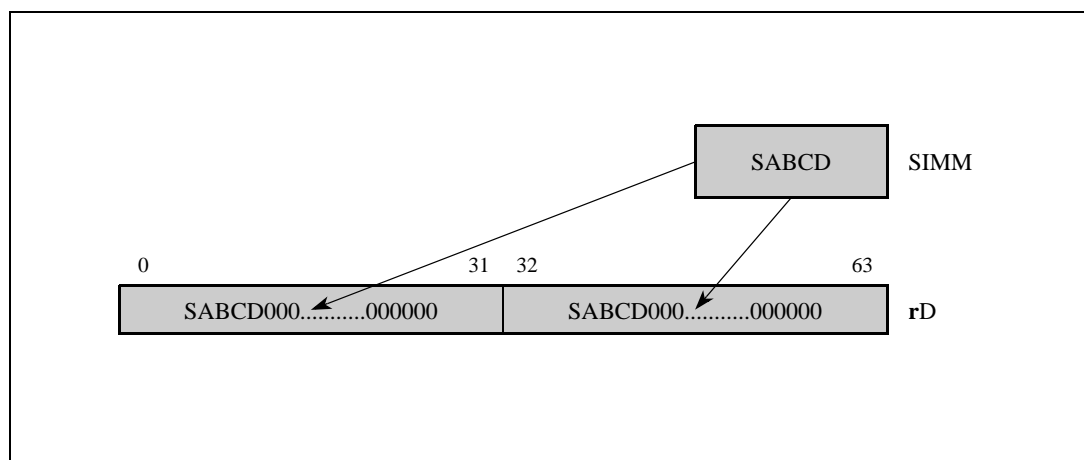


$$rD_{0:31} \leftarrow SIMM \parallel 27_0$$

$$rD_{32:63} \leftarrow SIMM \parallel 27_0$$

The 5-bit immediate value is padded with trailing zeros and placed in both elements of rD, as shown in [Figure 215](#). The SIMM ends up in bit positions rD[0–4] and rD[32–36].

**Figure 215. Vector splat fractional immediate (evsplatfi)**




**evsplatfi**

SPE APU	User
---------	------

**evsplatfi**

**Vector splat immediate**

**evsplatfi**                      rD,SIMM




$$rD_{0:31} \leftarrow \text{EXTS}(\text{SIMM})$$

$$rD_{32:63} \leftarrow \text{EXTS}(\text{SIMM})$$

The 5-bit immediate value is sign extended and placed in both elements of rD, as shown in [Figure 216](#).

Figure 216. evsrwsi sign extend




**evsrwis** **evsrwis**

SPE	APU	User
-----	-----	------

**Vector shift right word immediate signed**

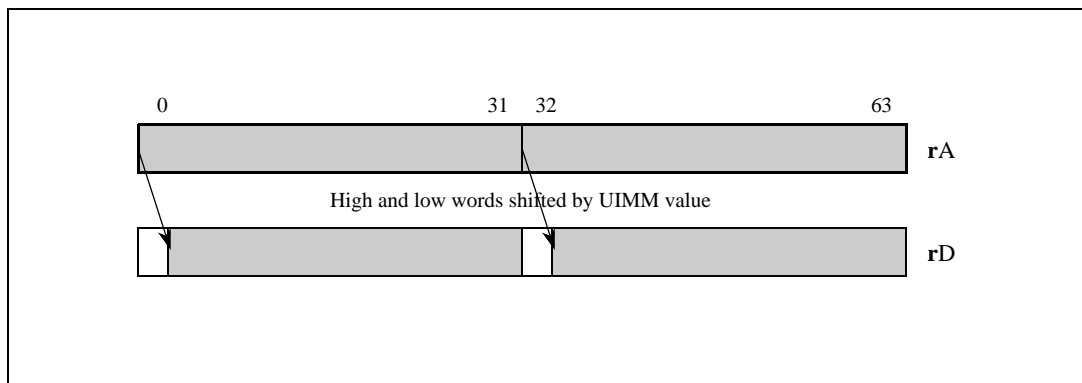
**evsrwis** **rD,rA,UIMM**



$n \leftarrow \text{UIMM}$   
 $rD_{0:31} \leftarrow \text{EXTS}(rA_{0:31-n})$   
 $rD_{32:63} \leftarrow \text{EXTS}(rA_{32:63-n})$

Both high and low elements of rA are shifted right by the 5-bit UIMM value. Bits in the most significant positions vacated by the shift are filled with a copy of the sign bit.

Figure 217. Vector shift right word immediate signed (evsrwiu)




**evsrwiu** **evsrwiu**

SPE	APU	User
-----	-----	------

**Vector shift right word immediate unsigned**

**evsrwiu** **rD,rA,UIMM**




```

n ← UIMM
rD0:31 ← EXTZ(rA0:31-n)
rD32:63 ← EXTZ(rA32:63-n)
    
```

Both high and low elements of rA are shifted right by the 5-bit UIMM value; 0 bits are shifted in to the most significant position. Bits in the most significant positions vacated by the shift are filled with a zero bit.


**Figure 218. Vector shift right word immediate unsigned (evsrwiu)**



**evsrws** SPE APU | User **evsrws**

**Vector shift right word signed**

**evsrws** rD,rA,rB



```

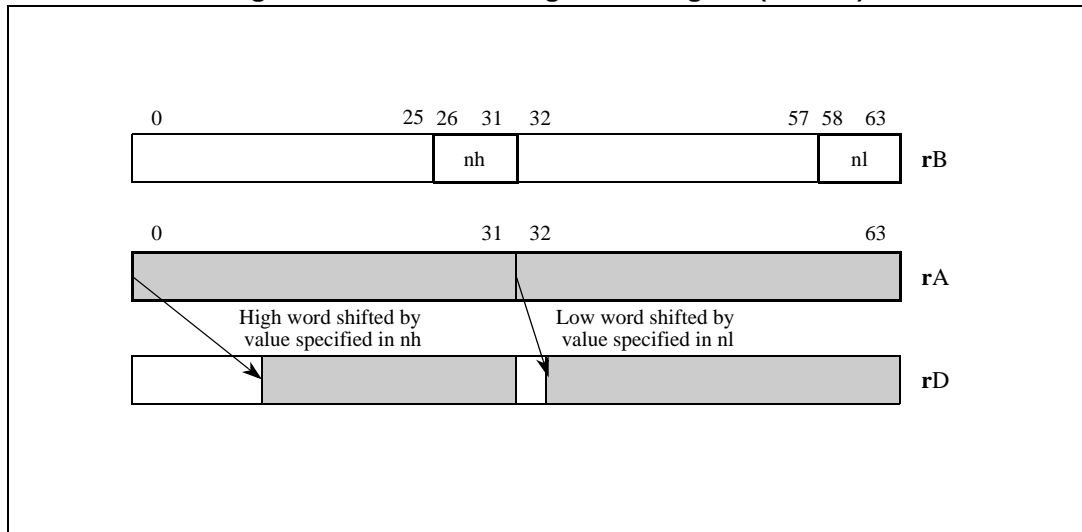
nh ← rB26:31
nl ← rB58:63
rD0:31 ← EXTS(rA0:31-nh)
rD32:63 ← EXTS(rA32:63-nl)
    
```

Both the high and low elements of rA are shifted right by an amount specified in rB. The result is placed into rD. 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 219. Vector shift right word signed (evsrws)



evsrwu
SPE APU
User
evsrwu  
**Vector shift right word unsigned**  
evsrwu
rD,rA,rB

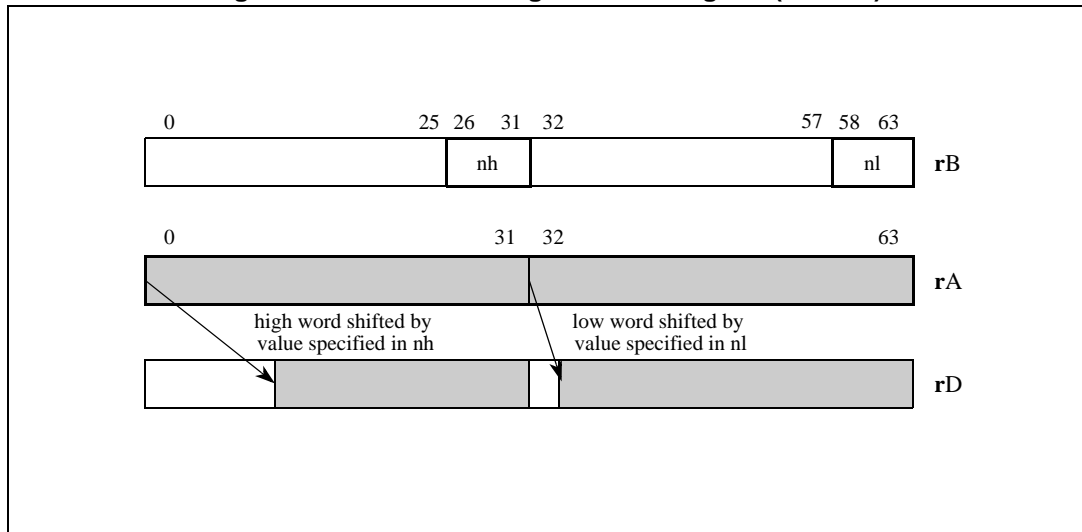
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	1	0	0	0	0	0	0

$$\begin{aligned}
 nh &\leftarrow rB_{26:31} \\
 nl &\leftarrow rB_{58:63} \\
 rD_{0:31} &\leftarrow EXTZ(rA_{0:31-nh}) \\
 rD_{32:63} &\leftarrow EXTZ(rA_{32:63-nl})
 \end{aligned}$$

Both the high and low elements of rA are shifted right by an amount specified in rB. The result is placed into rD. The separate shift amounts for each element are specified by 6 bits in rB 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 220. Vector shift right word unsigned (evsrwu)




**evstd**

SPE, Vector SPFP, Scalar DFPF APU	User
-----------------------------------	------

**evstd**

**Vector store double of double**

**evstd** **rS,d(rA)**



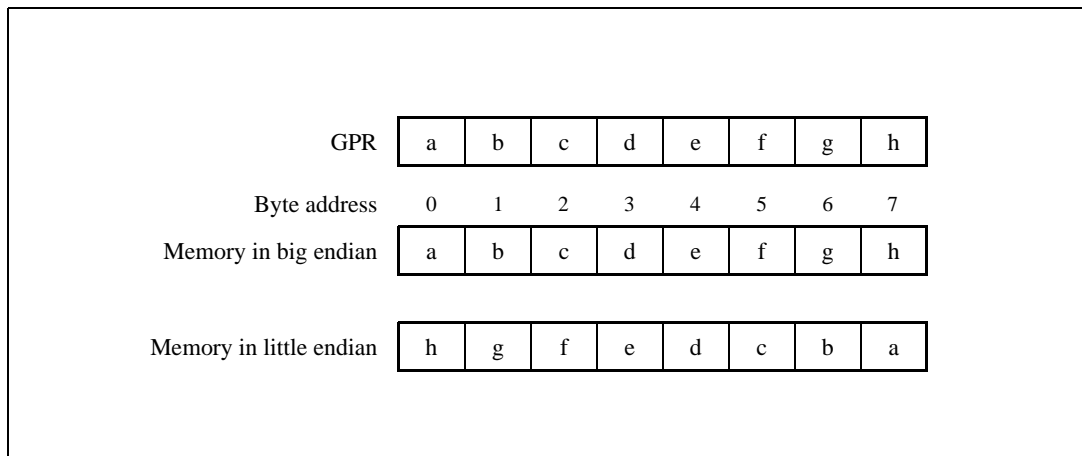
1.  $d = UIMM * 8$

if (rA = 0) then  $b \leftarrow 0$   
 else  $b \leftarrow (rA)$   
 $EA \leftarrow b + EXTZ(UIMM*8)$   
 $MEM(EA,8) \leftarrow RS_{0:63}$

The contents of rS are stored as a double word in storage addressed by EA.


Figure 221 shows how bytes are stored in memory as determined by the endian mode.

Figure 221. evstd results in big- and little-endian modes



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

**evstddx** SPE, Vector SPFP, Scalar DPFP APUs User **evstddx**  
**Vector store double of double indexed**  
**evstddx** **rS,rA,rB**

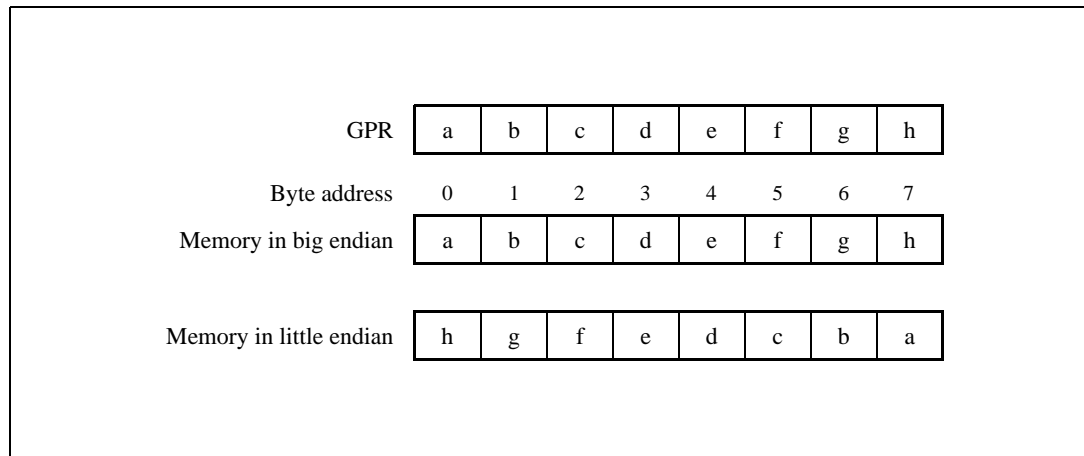


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 MEM(EA,8) ← RS<sub>0:63</sub>

The contents of rS are stored as a double word in storage addressed by EA.


Figure 222 shows how bytes are stored in memory as determined by the endian mode.

**Figure 222. evstddx Results in big- and little-endian modes**



*Note:* Implementation:  
 If the EA is not double-word aligned, an alignment exception occurs.

**evstdh** SPE APU User **evstdh**  
**Vector store double of four half words**  
**evstdh** **rS,d(rA)**



1. d = UIMM \* 8

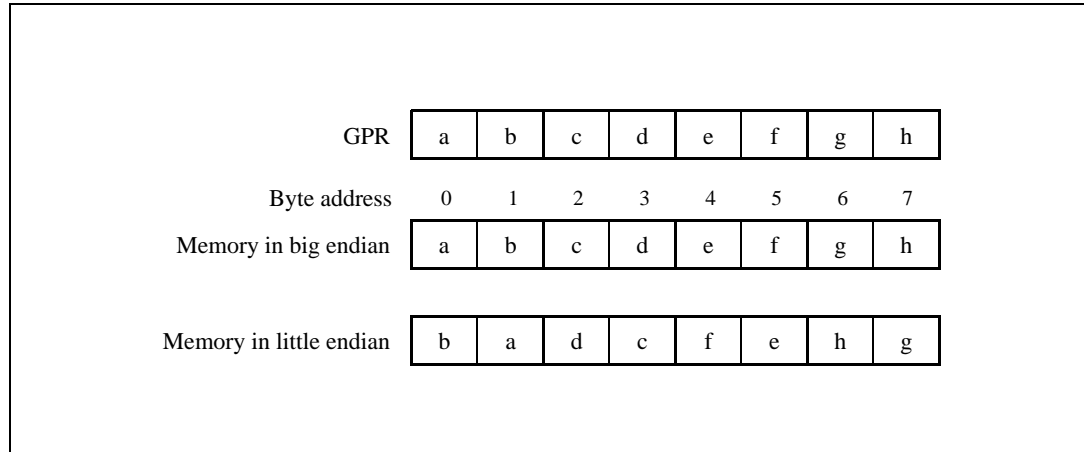
if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*8)  
 MEM(EA,2) ← RS<sub>0:15</sub>

MEM(EA+2,2) ← RS<sub>16:31</sub>  
 MEM(EA+4,2) ← RS<sub>32:47</sub>  
 MEM(EA+6,2) ← RS<sub>48:63</sub>

The contents of rS are stored as four half words in storage addressed by EA.

Figure 223 shows how bytes are stored in memory as determined by the endian mode.

Figure 223. evstdhx Results in big- and little-endian modes




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

**evstdhx** SPE APU User **evstdhx**

**Vector store double of four half words indexed**

**evstdhx** rS,rA,rB

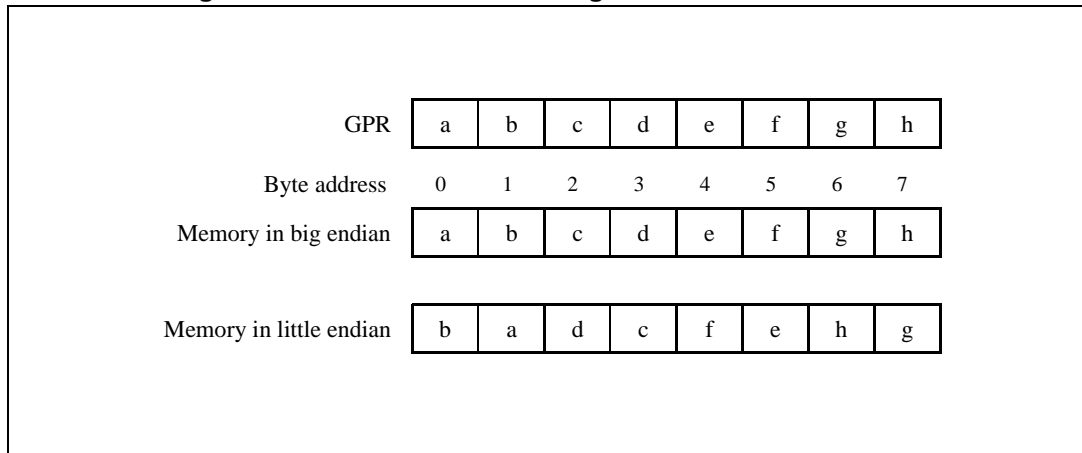


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 MEM(EA,2) ← RS<sub>0:15</sub>  
 MEM(EA+2,2) ← RS<sub>16:31</sub>  
 MEM(EA+4,2) ← RS<sub>32:47</sub>  
 MEM(EA+6,2) ← RS<sub>48:63</sub>



The contents of rS are stored as four half words in storage addressed by EA.

Figure 224 shows how bytes are stored in memory as determined by the endian mode.

Figure 224. evstdhx Results in big- and little-endian modes



Note: *Implementation:*  
 If the EA is not double-word aligned, an alignment exception occurs.



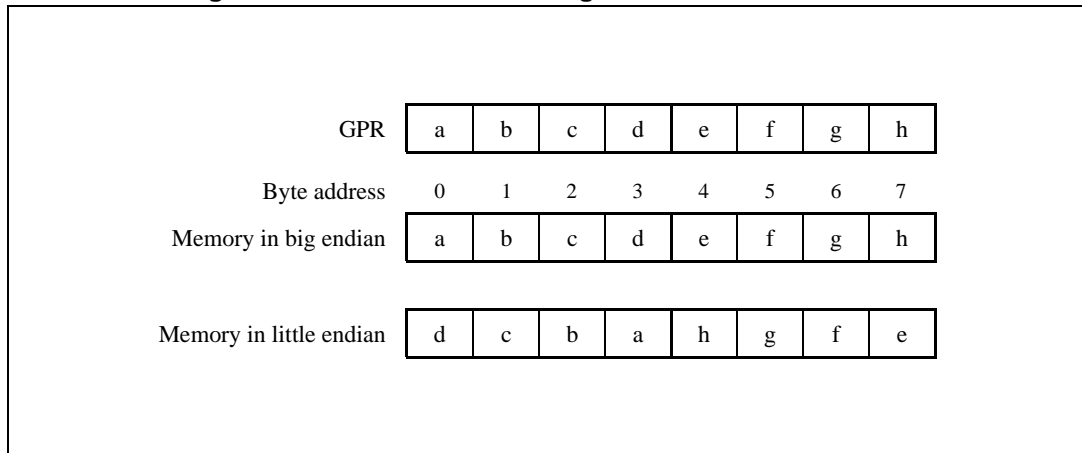
1. d = UIMM \* 8

if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*8)  
 MEM(EA,4) ← RS<sub>0:31</sub>  
 MEM(EA+4,4) ← RS<sub>32:63</sub>

The contents of rS are stored as two words in storage addressed by EA.

Figure 225 shows how bytes are stored in memory as determined by the endian mode.

Figure 225. evstdw results in big- and little-endian modes




Note: *Implementation:*  
 If the EA is not double-word aligned, an alignment exception occurs.

**evstdwx** SPE APU User **evstdwx**

**Vector store double of two words indexed**

**evstdwx** rS,rA,rB

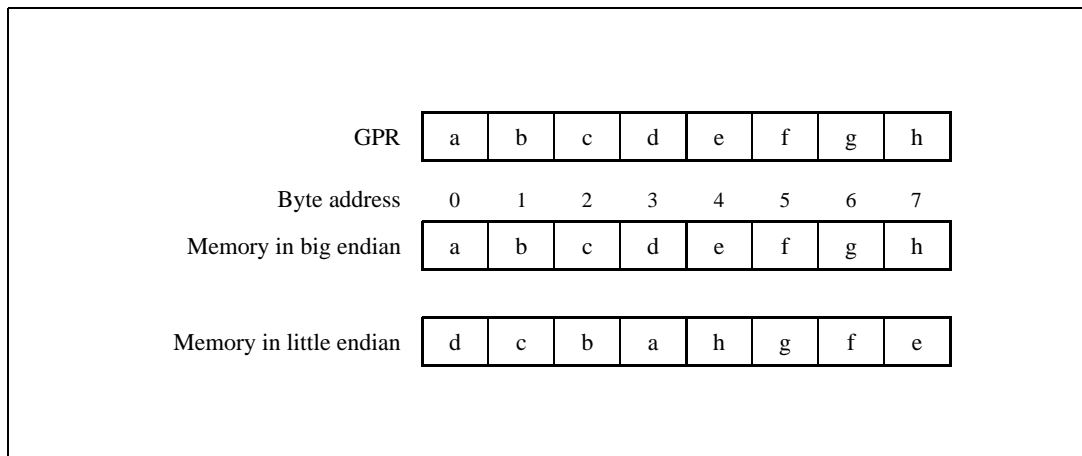


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 MEM(EA,4) ← RS<sub>0:31</sub>  
 MEM(EA+4,4) ← RS<sub>32:63</sub>

The contents of rS are stored as two words in storage addressed by EA.

Figure 226 shows how bytes are stored in memory as determined by the endian mode.

Figure 226. evstdwx Results in big- and little-endian modes



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

**evstwhe** **evstwhe**

SPE APU	User
---------	------

Vector store word of two half words from even

**evstwhe** **rS,d(rA)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rS						rA						UIMM <sup>(1)</sup>				0	1	1	0	0	1	1	0	0	0	1

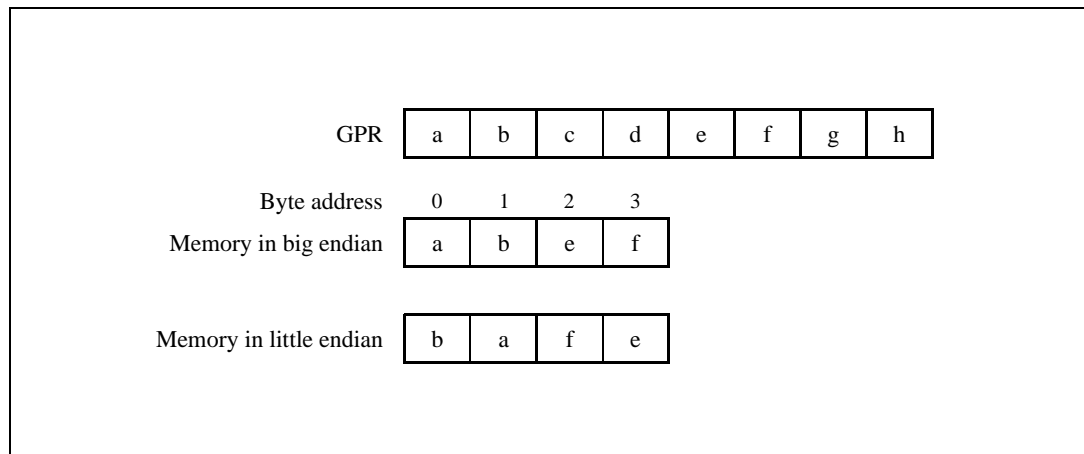
1.  $d = UIMM * 4$

if (rA = 0) then  $b \leftarrow 0$   
 else  $b \leftarrow (rA)$   
 $EA \leftarrow b + EXTZ(UIMM*4)$   
 $MEM(EA,2) \leftarrow RS_{0:15}$   
 $MEM(EA+2,2) \leftarrow RS_{32:47}$

The even half words from each element of rS are stored as two half words in storage addressed by EA.

Figure 227 shows how bytes are stored in memory as determined by the endian mode.

**Figure 227. evstwhe Results in big- and little-endian modes**



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

**evstwhex** **evstwhex**

SPE APU	User
---------	------

Vector store word of two half words from even indexed

**evstwhex** **rS,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	0	0	1	0	0	rS						rA						rB				0	1	1	0	0	1	1	0	0	0	0

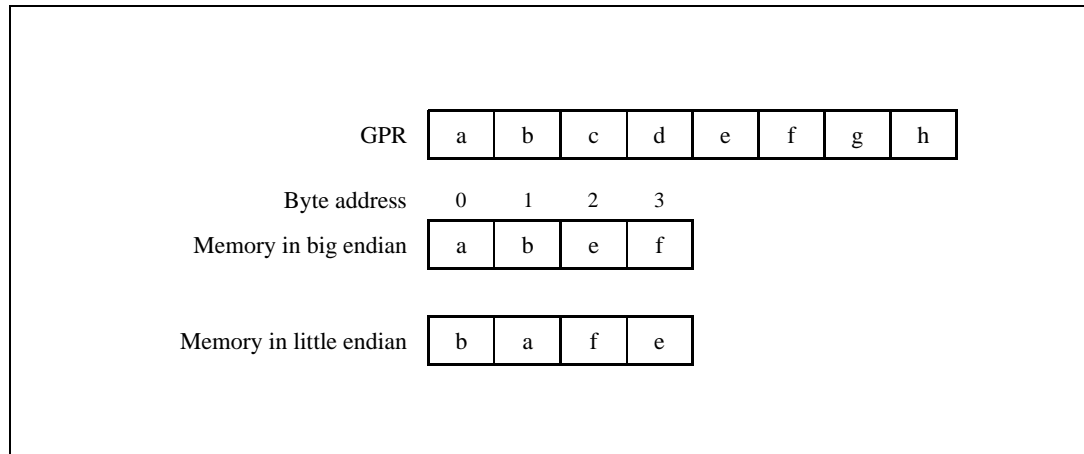
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + (rB)
MEM(EA,2) ← RS0:15
MEM(EA+2,2) ← RS32:47
    
```

The even half words from each element of rS are stored as two half words in storage addressed by EA.

Figure 228 shows how bytes are stored in memory as determined by the endian mode.

**Figure 228. evstwhex Results in big- and little-endian modes**



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

**evstwho** SPE APU User **evstwho**

**Vector store word of two half words from odd**

**evstwho** **rS,d(rA)**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rS					rA					UIMM <sup>(1)</sup>				0	1	1	0	0	1	1	0	1	0	1	

1. **d** = UIMM \* 4

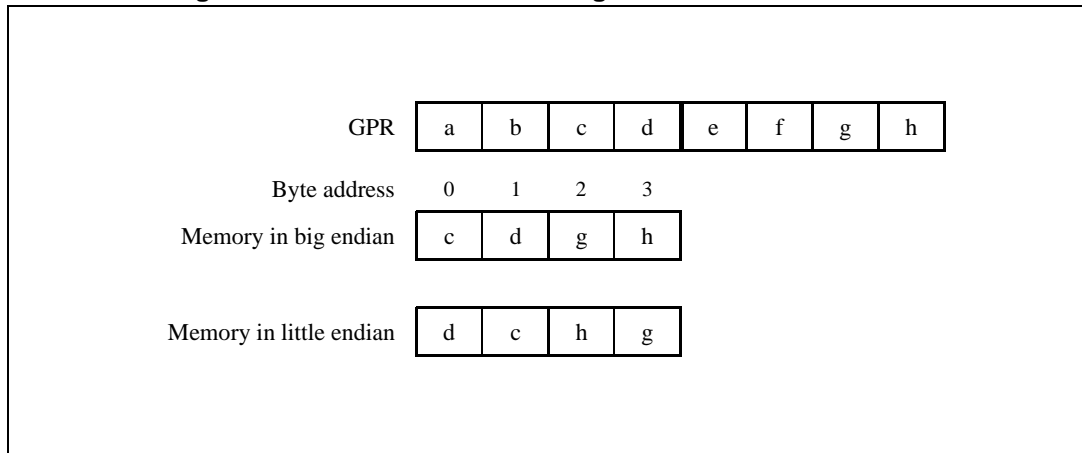
```

if (rA = 0) then b ← 0
else b ← (rA)
EA ← b + EXTZ(UIMM*4)
MEM(EA,2) ← RS16:31
MEM(EA+2,2) ← RS48:63
    
```


The odd half words from each element of rS are stored as two half words in storage addressed by EA.



**Figure 229. evstwho Results in big- and little-endian modes**



**evstwhox** SPE APU User **evstwhox**  
**Vector store word of two half words from odd indexed**  
**evstwhox** rS,rA,rB

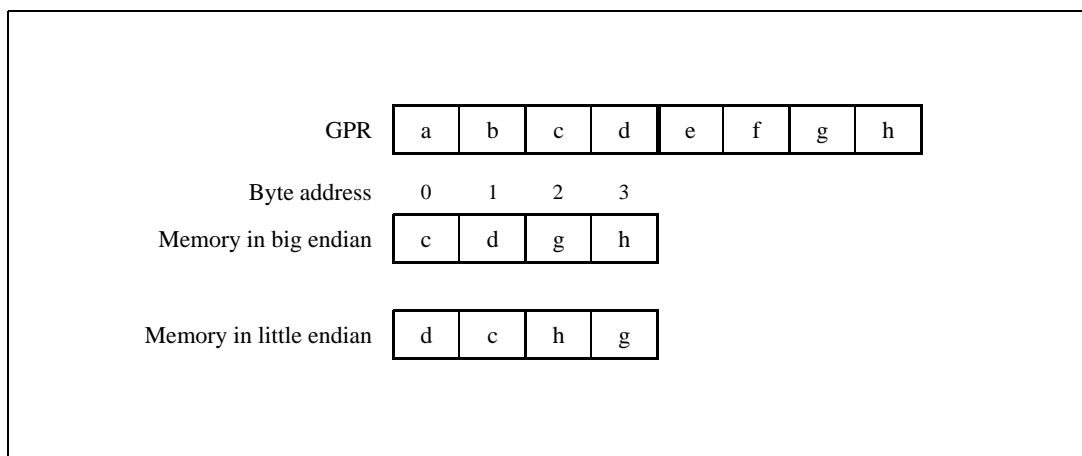


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 MEM(EA,2) ← RS<sub>16:31</sub>  
 MEM(EA+2,2) ← RS<sub>48:63</sub>

The odd half words from each element of rS are stored as two half words in storage addressed by EA.

Figure 230 shows how bytes are stored in memory as determined by the endian mode.

**Figure 230. evstwhox Results in big- and little-endian modes**




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

**evstwwe** SPE APU User **evstwwe**

Vector store word of word from even

**evstwwe** **rS,d(rA)**



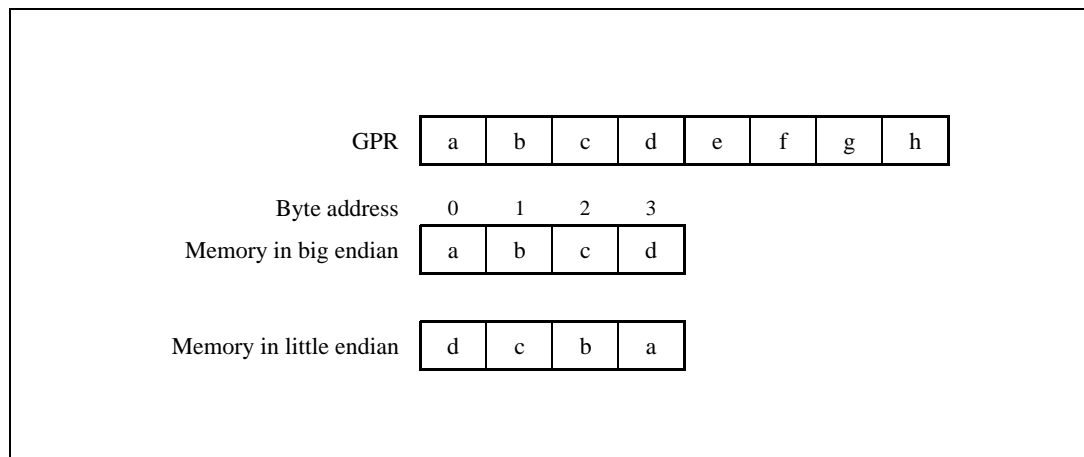
1.  $d = UIMM * 4$

if (rA = 0) then  $b \leftarrow 0$   
 else  $b \leftarrow (rA)$   
 $EA \leftarrow b + EXTZ(UIMM*4)$   
 $MEM(EA,4) \leftarrow RS_{0:31}$

The even word of rS is stored in storage addressed by EA.

Figure 231 shows how bytes are stored in memory as determined by the endian mode.

**Figure 231. evstwwe Results in big- and little-endian modes**




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

**evstwwex** SPE APU User **evstwwex**

Vector store word of word from even indexed

**evstwwex** **rS,rA,rB**



if (rA = 0) then  $b \leftarrow 0$   
 else  $b \leftarrow (rA)$

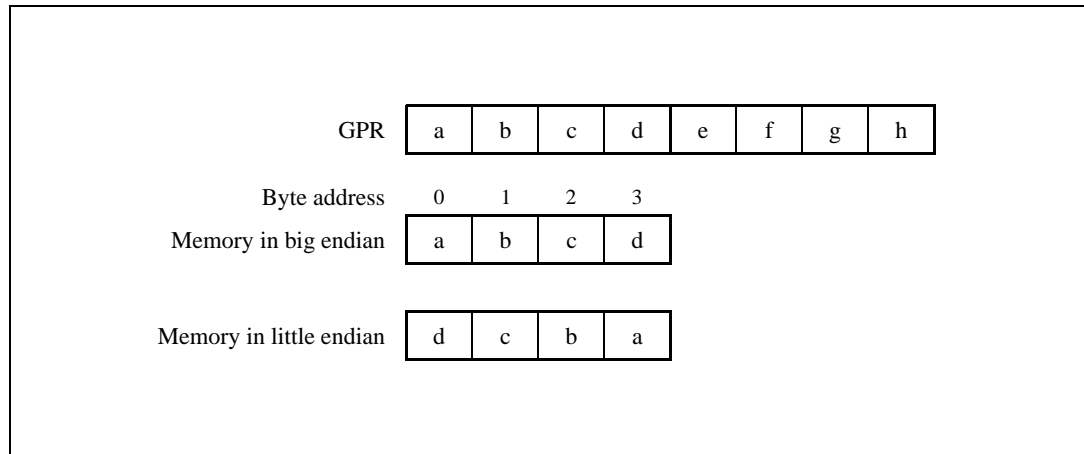
$$EA \leftarrow b + (rB)$$

$$MEM(EA,4) \leftarrow RS_{0:31}$$

The even word of rS is stored in storage addressed by EA.

Figure 232 shows how bytes are stored in memory as determined by the endian mode.

**Figure 232. evstwwex Results in big- and little-endian modes**




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

**evstwwo** **evstwwo**

SPE APU	User
---------	------

**Vector store word of word from odd**

**evstwwo** **rS,d(rA)**



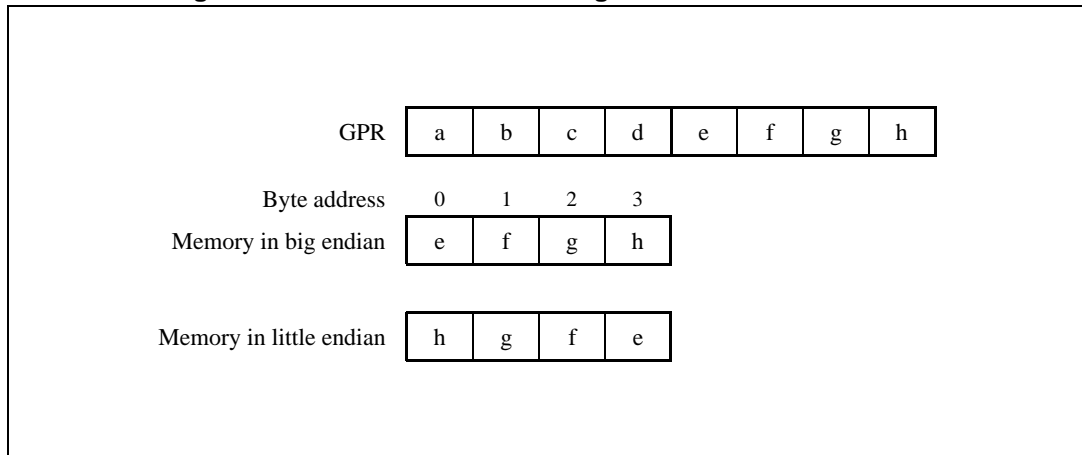
1. **d** = UIMM \* 4

if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + EXTZ(UIMM\*4)  
 MEM(EA,4) ← rS<sub>32:63</sub>

The odd word of rS is stored in storage addressed by EA.


Figure 233 shows how bytes are stored in memory as determined by the endian mode.

**Figure 233. evstwwo Results in big- and little-endian modes**



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

**evstwwox** SPE APU | User **evstwwox**  
**Vector store word of word from odd indexed**  
**evstwwox** **rS,rA,rB**

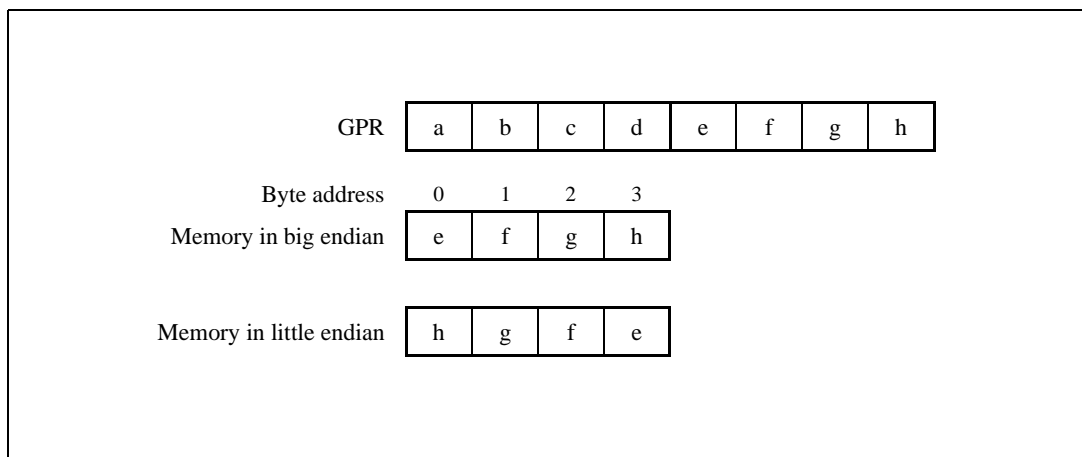


if (rA = 0) then b ← 0  
 else b ← (rA)  
 EA ← b + (rB)  
 MEM(EA,4) ← rS<sub>32:63</sub>

The odd word of rS is stored in storage addressed by EA.

Figure 234 shows how bytes are stored in memory as determined by the endian mode.

**Figure 234. evstwwox Results in big- and little-endian modes**



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

**evsubfsmiaaw**

SPE APU	User
---------	------

**evsubfsmiaaw**

**Vector subtract signed, modulo, integer to accumulator word**

**evsubfsmiaaw**                      **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						0	0	0	0	0	0	1	0	0	1	1	0	0	1	0	1	1

// high  
 $rD_{0:31} \leftarrow ACC_{0:31} - rA_{0:31}$


// low  
 $rD_{32:63} \leftarrow ACC_{32:63} - rA_{32:63}$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

Each word element in **rA** is subtracted from the corresponding element in the accumulator and the difference is placed into the corresponding **rD** word and into the accumulator.

Other registers altered: ACC

**Figure 235. Vector subtract signed, modulo, integer to accumulator word (evsubfsmiaaw)**




**evsubfssiaaw**

SPE APU	User
---------	------

**evsubfssiaaw**

**Vector subtract signed, saturate, integer to accumulator word**

**evsubfssiaaw**                      **rD,rA**



```

// high
temp0:63 ← EXTS(ACC0:31) - EXTS(rA0:31)
ovh ← temp31 ⊕ temp32
rD0:31 ← SATURATE(ovh, temp31, 0x80000000, 0x7ffffff, temp32:63)

// low
temp0:63 ← EXTS(ACC32:63) - EXTS(rA32:63)
ovl ← temp31 ⊕ temp32
rD32:63 ← SATURATE(ovl, temp31, 0x80000000, 0x7ffffff, temp32:63)


// update accumulator
ACC0:63 ← rD0:63

SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

Each signed integer word element in **rA** is sign-extended and subtracted from the corresponding sign-extended element in the accumulator, 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 236. Vector subtract signed, saturate, integer to accumulator word (evsubfssiaaw)**



**evsubfumiaaw**

SPE APU	User
---------	------

**evsubfumiaaw**

**Vector subtract unsigned, modulo, integer to accumulator word**

**evsubfumiaaw**                      rD,rA

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					0	0	0	0	0	1	0	0	1	1	0	0	1	0	1	0

// high  
 $rD_{0:31} \leftarrow ACC_{0:31} - rA_{0:31}$


// low  
 $rD_{32:63} \leftarrow ACC_{32:63} - rA_{32:63}$

// update accumulator  
 $ACC_{0:63} \leftarrow rD_{0:63}$

Each unsigned integer word element in rA is subtracted from the corresponding element in the accumulator and the results are placed in rD and into the accumulator.

Other registers altered: ACC

**Figure 237. Vector subtract unsigned, modulo, integer to accumulator word (evsubfumiaaw)**



**evsubfusiaaw**

SPE APU	User
---------	------

**evsubfusiaaw**

**Vector subtract unsigned, saturate, integer to accumulator word**

**evsubfusiaaw**                      **rD,rA**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	0	1	0	0	rD						rA						0	0	0	0	0	0	1	0	0	1	1	0	0	0	0	1	0

```

// high
temp0:63 ← EXTZ(ACC0:31) - EXTZ(rA0:31)
ovh ← temp31
rD0:31 ← SATURATE(ovh, temp31, 0x00000000, 0x00000000, temp32:63)

// low
temp0:63 ← EXTS(ACC32:63) - EXTS(rA32:63)
ovl ← temp31
rD32:63 ← SATURATE(ovl, temp31, 0x00000000, 0x00000000, temp32:63)


// update accumulator
ACC0:63 ← rD0:63

SPEFSCROVH ← ovh
SPEFSCROV ← ovl
SPEFSCRSOVH ← SPEFSCRSOVH | ovh
SPEFSCRSOV ← SPEFSCRSOV | ovl
    
```

Each unsigned integer word element in rA is zero-extended and subtracted from the corresponding zero-extended element in the accumulator, saturating if underflow occurs, and the results are placed in rD and the accumulator. Any underflow is recorded in the SPEFSCR overflow and summary overflow bits.

Other registers altered: SPEFSCR ACC

**Figure 238. Vector subtract unsigned, saturate, integer to accumulator word (evsubfusiaaw)**





**evsubfw**

SPE APU	User
---------	------

**evsubfw**

**Vector subtract from word**

**evsubfw**                      **rD,rA,rB**


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31					
0	0	0	1	0	0	rD						rA						rB						0	1	0	0	0	0	0	0	0	0	1	0	0

$$rD_{0:31} \leftarrow rB_{0:31} - rA_{0:31} \quad // \text{ Modulo difference}$$

$$rD_{32:63} \leftarrow rB_{32:63} - rA_{32:63} \quad // \text{ Modulo difference}$$

Each signed integer element of **rA** is subtracted from the corresponding element of **rB** and the results are placed into **rD**.

**Figure 239. Vector subtract from word (evsubfw)**



**evsubifw**

SPE APU	User
---------	------

**evsubifw**

**Vector subtract immediate from word**

**evsubifw**                      **rD,UIMM,rB**


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	0	0	1	0	0	rD						UIMM						rB						0	1	0	0	0	0	0	0	0	1	1	0

$$rD_{0:31} \leftarrow rB_{0:31} - \text{EXTZ}(UIMM) \quad // \text{ Modulo difference}$$

$$rD_{32:63} \leftarrow rB_{32:63} - \text{EXTZ}(UIMM) \quad // \text{ Modulo difference}$$

**UIMM** is zero-extended and subtracted from both the high and low elements of **rB**. Note that the same value is subtracted from both elements of the register. **UIMM** is 5 bits.

Figure 240. Vector subtract immediate from word (evsubifw)



**evxor** SPE APU User **evxor**

**Vector XOR**

**evxor** rD,rA,rB

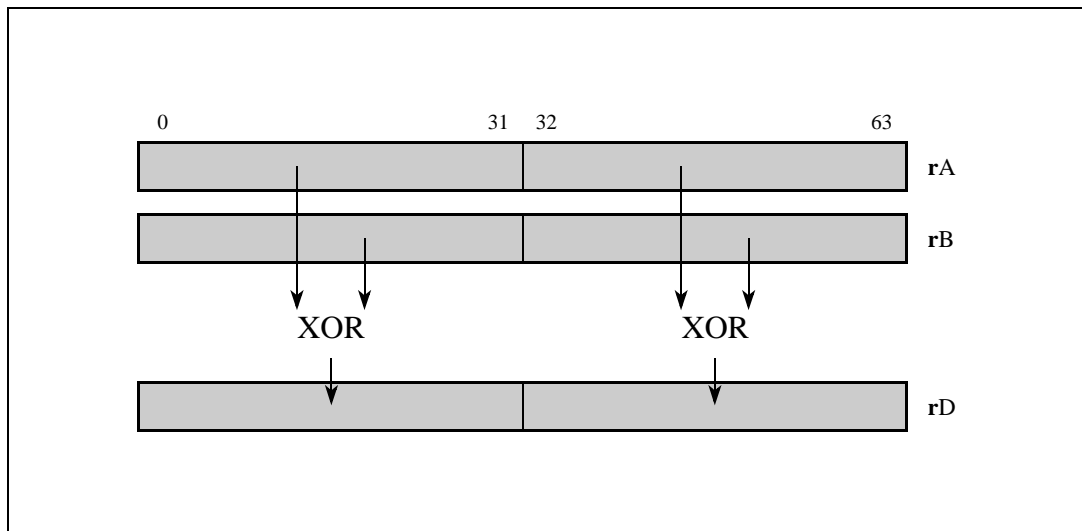
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	0	0	rD					rA					rB					0	1	0	0	0	0	1	0	1	1	0

$$rD_{0:31} \leftarrow rA_{0:31} \oplus rB_{0:31} \quad // \text{ Bitwise XOR}$$


$$rD_{32:63} \leftarrow rA_{32:63} \oplus rB_{32:63} \quad // \text{ Bitwise XOR}$$

Each element of rA and rB is exclusive-ORed. The results are placed in rD.

Figure 241. Vector XOR (evxor)



<b>extsb</b>		Book E	User		<b>extsb</b>
<b>Extend sign (byte   half word)</b>					
<b>extsb</b>	<b>rA,rS</b>				(SZ=0b01, Rc=0)
<b>extsb.</b>	<b>rA,rS</b>				(SZ=0b01, Rc=1)
<b>extsh</b>	<b>rA,rS</b>				(SZ=0b00, Rc=0)
<b>extsh.</b>	<b>rA,rS</b>				(SZ=0b00, Rc=1)



```

if 'extsb[.]' then n ← 56
if 'extsh[.]' then n ← 48
if 'extsw' then n ← 32
if Rc=1 then do
    LT ← rSn:63 < 0
    GT ← rSn:63 > 0
    EQ ← rSn:63 = 0
    CR0 ← LT || GT || EQ || SO



s ← rSn
rA ← ns || rSn:63
    
```

For **extsb[.]**, the contents of rS[56–63] are placed into rA[56–63]. Bit rS[56] is copied into bits 0–55 of rA. If Rc=1, CR field 0 is set to reflect the result.

For **extsh[.]**, the contents of rS[48–63] are placed into rA[48–63]. rS[48] is copied into rA[0–47]. If Rc=1, CR field 0 is set to reflect the result.

- Other registers altered:  
CR0 (if Rc=1)

<b>_extsbx</b>		VLE	User		<b>_extsbx</b>
<b>Extend Sign (Byte   Halfword)</b>					



```

if se_extsb then n ← 56
if se_extsh then n ← 48
if 'extsw' then n ← 32
if Rc=1 then do
    LT ← GPR(RS)n:63 < 0
    GT ← GPR(RS)n:63 > 0
    EQ ← GPR(RS)n:63 = 0
    CR0 ← LT || GT || EQ || SO
s ← GPR(RS or RX)n
    
```



$$GPR(RA \text{ or } RX) \leftarrow {}^{n-32}s \parallel GPR(RS \text{ or } RX)_{n:63}$$

For **se\_extsb**, the contents of bits 56–63 of GPR(rX) are placed into bits 56–63 of GPR(rX). Bit 56 of the contents of GPR(rX) is copied into bits 32–55 of GPR(rX).

For **se\_extsh**, the contents of bits 48–63 of GPR(rX) are placed into bits 48–63 of GPR(rX). Bit 48 of the contents of GPR(rX) is copied into bits 32–47 of GPR(rX).



Special Registers Altered: CR0 (if Rc=1)

**\_extzx**

VLE	User
-----	------

**\_extzx**

**Extend Zero (Byte | Halfword)**



if 'se\_extzb' then n ← 56  
 if 'se\_extzh' then n ← 48  
 $GPR(RX) \leftarrow {}^{n-32}0 \parallel GPR(RX)_{n:63}$

For **se\_extzb**, the contents of bits 56–63 of GPR(rX) are placed into bits 56–63 of GPR(rX). Bits 32–55 of GPR(rX) are cleared.

For **se\_extzh**, the contents of bits 48–63 of GPR(rX) are placed into bits 48–63 of GPR(rX). Bits 32–47 of GPR(rX) are cleared.

Special Registers Altered: None


**fabs**

Book E	User
--------	------

**fabs**

**Floating absolute value**

**fabs** **frD,frB** (Rc=0)  
**fabs.** **frD,frB** (Rc=1)



$$frD \leftarrow 0b0 \parallel frB_{1:63}$$

The contents of frB with bit 0 cleared are placed into frD.

If MSR[FP]=0, an attempt to execute **fabs[.]** causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)


**fadd**

Book E	User
--------	------

**fadd**

**Floating add [single]**

<b>fadd</b>	<b>frD,frA,frB</b>	(P=1, Rc=0)
<b>fadd.</b>	<b>frD,frA,frB</b>	(P=1, Rc=1)
<b>fadds</b>	<b>frD,frA,frB</b>	(P=0, Rc=0)
<b>fadds.</b>	<b>frD,frA,frB</b>	(P=0, Rc=1)



if P=1 then frD ← frA +<sub>dp</sub> frB  
 else frD ← frA +<sub>sp</sub> frB

The floating-point operand in **frA** is added to the floating-point operand in **frB**.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into **frD**.

Floating-point addition is based on exponent comparison and addition of the two significands. The exponents of the two operands are compared, and the significand accompanying the smaller exponent is shifted right, with its exponent increased by one for each bit shifted, until the two exponents are equal. The two significands are then added or subtracted as appropriate, depending on the signs of the operands, to form an intermediate sum. All 53 bits of the significand as well as all three guard bits (G, R, and X) enter into the computation.

If a carry occurs, the sum’s significand is shifted right one bit position and the exponent is increased by one.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fadd[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN VXISI  
 CR1 ← FX || FEX || VX || OX (if Rc=1)


**fcfid**

Book E	User
--------	------

**fcfid**

**Floating convert from integer doubleword**

**fcfid** frD,frB



sign ← frB<sub>0</sub>  
 exp ← 63  
 frac<sub>0:63</sub> ← frB  
 If frac<sub>0:63</sub> = 0 then go to Zero Operand  
 If sign = 1 then frac<sub>0:63</sub> ← -frac<sub>0:63</sub> + 1



```

Do while frac0 = 0 /* do loop 0 times if frB = max negative integer */
    frac0:63 ← frac1:63 || 0b0
    exp ← exp : 1
End
Round Float( sign, exp, frac0:63, FPSCR[RN] )
If sign = 0 then FPSCR[FPRF] ← '+normal number'
If sign = 1 then FPSCR[FPRF] ← ':normal number'
frD0 ← sign
frD[1-11] ← exp + 1023 /* exp + bias */
frD[12-63] ← frac1:52
Done

Zero Operand:
FPSCR[FR,FI] ← 0b00
FPSCR[FPRF] ← '+zero'
frD ← 0x0000_0000_0000_0000
Done

Round Float( sign, exp, frac0:63, round_mode ):
inc ← 0
lsb ← frac52
gbit ← frac53
rbit ← frac54
xbit ← frac55:63 > 0
If round_mode = 0b00 then
    Do /* comparison ignores u bits */
        If sign || lsb || gbit || rbit || xbit = 0bu11uu then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0bu011u then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0bu01u1 then
            inc ← 1
    End
If round_mode = 0b10 then
    Do /* comparison ignores u bits */
        If sign || lsb || gbit || rbit || xbit = 0b0u1uu then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0b0uu1u then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0b0uuu1 then
            inc ← 1
    End
If round_mode = 0b11 then
    Do /* comparison ignores u bits */
        If sign || lsb || gbit || rbit || xbit = 0b1u1uu then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0b1uu1u then
            inc ← 1
        If sign || lsb || gbit || rbit || xbit = 0b1uuu1 then
            inc ← 1
    End
frac0:52 ← frac0:52 + inc
If carry_out = 1 then exp ← exp + 1
FPSCR[FR] ← inc

```

FPSCR[FI] ← gbit | rbit | xbit  
 FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]  
 Return

The 64-bit signed operand in **frB** is converted to an infinitely precise floating-point integer. The result of the conversion is rounded to double-precision, as specified by FPSCR[RN], and placed into **frD**.

FPSCR[FPRF] is set to the class and sign of the result. FPSCR[FR] is set if the result is incremented when rounded. FPSCR[FI] is set if the result is inexact.

If MSR[FP]=0, an attempt to execute **fcfid** causes a floating-point unavailable interrupt.

Other registers altered: FPRF FR FI FX XX


**fcmpu**

Book E	User
--------	------

**fcmpu**

**Floating compare**

**fcmpu** crD,frA,frB (U=0)  
**fcmpo** crD,frA,frB (U=1)



```

if frA is a NaN or
frB is a NaN then    c ← 0b0001
else if frA < frB then c ← 0b1000
else if frA > frB then c ← 0b0100
else                  c ← 0b0010
FPCC ← c
CR4×crD:4×crD+3 ← c
if 'fcmpu' & (frA is a SNaN or frB is a SNaN) then VXSNaN ← 1
if 'fcmpo' then do
    if frA is a SNaN or frB is a SNaN then do
        if VE=0 then VXVC ← 1
    else if frA is a QNaN or frB is a QNaN then VXVC ← 1
    
```

The floating-point operand in **frA** is compared to the floating-point operand in **frB**. The result of the compare is placed into CR field **crD** and the FPCC.

If either of the operands is a NaN, either quiet or signaling, the CR field **crD** and the FPCC are set to reflect unordered.

If **fcmpu**, then if either of the operands is a signaling NaN, VXSNaN is set.

If **fcmpo**, then do the following:

If either of the operands is a signaling NaN and invalid operation is disabled (VE=0), VXVC is set. If neither operand is a signaling NaN but at least one operand is a quiet NaN, then VXVC is set.

If MSR[FP]=0, an attempt to execute **fcmpo** or **fcmpu** causes a floating-point unavailable interrupt.



Other registers altered:

- CR field crD  
FPCC FX VXSNaN  
VXVC(if **fcmpo**)

<b>fctid</b>	Book E	User	<b>fctid</b>
<b>Floating convert to integer doubleword</b>			
<b>fctid</b>	<b>frD,frB</b>		(Z=0)
<b>fctidz</b>	<b>frD,frB</b>		(Z=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
1	1	1	1	1	1	<b>frD</b>						///						<b>frB</b>						1	1	0	0	1	0	1	1	1	Z	/

```

if 'fctid[.]' then round_mode ← FPSCR[RN]
if 'fctidz[.]' then round_mode ← 0b01
sign ← frB0
If frB[1:11] = 2047 and frB[12:63] = 0 then goto Infinity Operand
If frB[1:11] = 2047 and frB12 = 0 then goto SNaN Operand
If frB[1:11] = 2047 and frB12 = 1 then goto QNaN Operand
If frB[1:11] > 1086 then goto Large Operand
If frB[1:11] > 0 then exp ← frB[1:11] : 1023 /* exp : bias */
If frB[1:11] = 0 then exp ← :1022
/* normal; need leading 0 for later complement */
If frB[1:11] > 0 then frac0:64 ← 0b01 || frB[12:63] || 110
/* denormal */
If frB[1:11] = 0 then frac0:64 ← 0b00 || frB[12:63] || 110
gbit || rbit || xbit ← 0b000
Do i=1,63:exp /* do the loop 0 times if exp = 63 */
    frac0:64 || gbit || rbit || xbit ← 0b0 || frac0:64 || gbit || (rbit |
xbit)
End
Round Integer( sign, frac0:64, gbit, rbit, xbit, round_mode )
/* needed leading 0 for :264 < frB < :263 */
If sign=1 then frac0:64 ← -frac0:64 + 1
If frac0:64 > 263:1 then goto Large Operand
If frac0:64 < :263 then goto Large Operand
FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
FPSCR[FPRF] ← undefined
frD ← frac1:64
Done
Round Integer( sign,frac0:64, gbit, rbit, xbit, round_mode ):
inc ← 0
If round_mode = 0b00 then /* comparison ignores u bits */
    Do
        If sign || frac64 || gbit || rbit || xbit = 0bu11uu
            then inc ← 1
        If sign || frac64 || gbit || rbit || xbit = 0bu011u
            then inc ← 1
        If sign || frac64 || gbit || rbit || xbit = 0bu01u1
            then inc ← 1
    
```





```

        End
        If round_mode = 0b10 then /* comparison ignores u bits */
        Do
            If sign || frac64 || gbit || rbit || xbit = 0b0u1uu
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b0uu1u
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b0uuu1
then inc ← 1
        End
        If round_mode = 0b11 then /* comparison ignores u bits */
        Do
            If sign || frac64 || gbit || rbit || xbit = 0b1u1uu
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b1uu1u
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b1uuu1
then inc ← 1
        End
        frac0:64 ← frac0:64 + inc
        FPSCR[FR] ← inc
        FPSCR[FI] ← gbit | rbit | xbit
        Return
Infinity Operand:
        FPSCR[FR,FI,VXCVI] ← 0b001
        If FPSCR[VE] = 0 then Do
            If sign = 0 then frD ← 0x7FFF_FFFF_FFFF_FFFF
            If sign = 1 then frD ← 0x8000_0000_0000_0000
            FPSCR[FPRF] ← undefined
        End
        Done
SNaN Operand:
        FPSCR[FR,FI,VXSNAN,VXCVI] ← 0b0011
        If FPSCR[VE] = 0 then Do
            frD ← 0x8000_0000_0000_0000
            FPSCR[FPRF] ← undefined
        End
        Done
QNaN Operand:
        FPSCR[FR,FI,VXCVI] ← 0b001
        If FPSCR[VE] = 0 then Do
            frD ← 0x8000_0000_0000_0000
            FPSCR[FPRF] ← undefined
        End
        Done
Large Operand:
        FPSCR[FR,FI,VXCVI] ← 0b001
        If FPSCR[VE] = 0 then Do
            If sign = 0 then frD ← 0x7FFF_FFFF_FFFF_FFFF
            If sign = 1 then frD ← 0x8000_0000_0000_0000
            FPSCR[FPRF] ← undefined

```

End  
Done

For **fctid** or **fctid.**, the rounding mode is specified by FPSCR[RN].

For **fctidz** or **fctidz.**, the rounding mode used is round toward zero.

The floating-point operand in **frB** is converted to a 64-bit signed integer, using the rounding mode specified by the instruction, and placed into **frD**.

If the floating-point operand in **frB** is greater than  $2^{63}-1$ , then 0x7FFF\_FFFF\_FFFF\_FFFF is placed into **frD**. If the floating-point operand in **frB** is less than  $-2^{63}$ , 0x8000\_0000\_0000\_0000 is placed into **frD**.

Except for enabled invalid operation exceptions, FPSCR[FPRF] is undefined. FPSCR[FR] is set if the result is incremented when rounded. FPSCR[FI] is set if the result is inexact.

If MSR[FP]=0, an attempt to execute **fctid[z]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF (undefined) FR FI FX XX VXSNaN VXCVI


**fctiw**

Book E	User
--------	------

**fctiw**

**Floating convert to integer word**

<b>fctiw</b>	<b>frD,frB</b>	(Z=0, Rc=0)
<b>fctiw.</b>	<b>frD,frB</b>	(Z=0, Rc=1)
<b>fctiwz</b>	<b>frD,frB</b>	(Z=1, Rc=0)
<b>fctiwz.</b>	<b>frD,frB</b>	(Z=1, Rc=1)



```

if 'fctiw[.]' then round_mode ← FPSCR[RN]
if 'fctiwz[.]' then round_mode ← 0b01
sign ← frB0
If frB[1:11] = 2047 and frB[12:63] = 0 then goto Infinity Operand
If frB[1:11] = 2047 and frB12 = 0 then goto SNaN Operand
If frB[1:11] = 2047 and frB12 = 1 then goto QNaN Operand
If frB[1:11] > 1086 then goto Large Operand
If frB[1:11] > 0 then exp ← frB[1:11] : 1023 /* exp : bias */
If frB[1:11] = 0 then exp ← :1022
/* normal; need leading 0 for later complement */
If frB[1:11] > 0 then frac0:64 ← 0b01 || frB[12:63] || 110
/* denormal */
If frB[1:11] = 0 then frac0:64 ← 0b00 || frB[12:63] || 110
gbit || rbit || xbit ← 0b000
Do i=1,63:exp /* do the loop 0 times if exp = 63 */
    frac0:64 || gbit || rbit || xbit ← 0b0 || frac0:64 || gbit || (rbit |
xbit)
End
Round Integer( sign, frac0:64, gbit, rbit, xbit, round_mode )
/* needed leading 0 for :264 < frB < :263 */
If sign=1 then frac0:64 ← -frac0:64 + 1
    
```



```

    If frac0:64 > 231:1 then goto Large Operand
    If frac0:64 < :231 then goto Large Operand
    FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
    frD ← 0xuuuu_uuuu || frac33:64 /* u is undefined hex digit */
    FPSCR[FPRF] ← undefined
    Done
Round Integer( sign, frac0:64, gbit, rbit, xbit, round_mode ):
    inc ← 0
    If round_mode = 0b00 then /* comparison ignores u bits */
        Do
            If sign || frac64 || gbit || rbit || xbit = 0bu11uu
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0bu011u
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0bu01u1
then inc ← 1
        End
    If round_mode = 0b10 then /* comparison ignores u bits */
        Do
            If sign || frac64 || gbit || rbit || xbit = 0b0u1uu
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b0uu1u
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b0uuu1
then inc ← 1
        End
    If round_mode = 0b11 then /* comparison ignores u bits */
        Do
            If sign || frac64 || gbit || rbit || xbit = 0b1u1uu
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b1uu1u
then inc ← 1
            If sign || frac64 || gbit || rbit || xbit = 0b1uuu1
then inc ← 1
        End
    End
    frac0:64 ← frac0:64 + inc
    FPSCR[FR] ← inc
    FPSCR[FI] ← gbit | rbit | xbit
    Return
Infinity Operand:
    FPSCR[FR,FI,VXCVI] ← 0b001
    If FPSCR[VE] = 0 then Do /* u is undefined hex digit */
        If sign = 0 then frD ← 0xuuuu_uuuu_7FFF_FFFF
        If sign = 1 then frD ← 0xuuuu_uuuu_8000_0000
        FPSCR[FPRF] ← undefined
    End
    Done
SNaN Operand:
    FPSCR[FR,FI,VXSNAN,VXCVI] ← 0b0011
    If FPSCR[VE] = 0 then Do /* u is undefined hex digit */
        frD ← 0xuuuu_uuuu_8000_0000

```

```

                                FPSCR[FPRF] ← undefined
                                End
                                Done
    QNaN Operand:
                                FPSCR[FR,FI,VXCVI] ← 0b001
                                If FPSCR[VE] = 0 then Do /* u is undefined hex digit */
                                    frD ← 0xuuuu_uuuu_8000_0000
                                    FPSCR[FPRF] ← undefined
                                End
                                Done
    Large Operand:
                                FPSCR[FR,FI,VXCVI] ← 0b001
                                If FPSCR[VE] = 0 then Do /* u is undefined hex digit */
                                    If sign = 0 then frD ← 0xuuuu_uuuu_7FFF_FFFF
                                    If sign = 1 then frD ← 0xuuuu_uuuu_8000_0000
                                    FPSCR[FPRF] ← undefined
                                End
                                Done
    
```

For **fctiw** or **fctiw.**, the rounding mode is specified by FPSCR[RN].

For **fctiwz** or **fctiwz.**, the rounding mode used is round toward zero.

The floating-point operand in **frB** is converted to a 32-bit signed integer, using the rounding mode specified by the instruction, and placed into **frD[32–63]**; **frD[0–31]** are undefined.

If the operand in **frB** is greater than  $2^{31}-1$ , then **frD[32–63]** are set to 0x7FFF\_FFFF. If the operand in **frB** is less than  $-2^{31}$ , then **frD[32–63]** are set to 0x8000\_0000.

Except for enabled invalid operation exceptions, FPSCR[FPRF] is undefined. FPSCR[FR] is set if the result is incremented when rounded. FPSCR[FI] is set if the result is inexact.

If MSR[FP]=0, an attempt to execute **fctiw[z].** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF (undefined) FR FI FX XX VXSNaN VXCVI  
CR1 ← FX || FEX || VX || OX (if Rc=1)


**fdiv**

Book E	User
--------	------

**fdiv**

**Floating divide [single]**

<b>fdiv</b>	<b>frD,frA,frB</b>	(P=1, Rc=0)
<b>fdiv.</b>	<b>frD,frA,frB</b>	(P=1, Rc=1)
<b>fdivs</b>	<b>frD,frA,frB</b>	(P=0, Rc=0)
<b>fdivs.</b>	<b>frD,frA,frB</b>	(P=0, Rc=1)



```

if P=1 then frD ← frA ÷dp frB
else frD ← frA ÷sp frB
    
```

The floating-point operand in **frA** is divided by the floating-point operand in **frB**. The remainder is not supplied as a result.



If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into **frD**.

Floating-point division is based on exponent subtraction and division of the significands.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1 and zero divide exceptions when FPSCR[ZE]=1.

If MSR[FP]=0, an attempt to execute **fdiv[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX ZX XX VXSAN VXIDI VXZDZ  
CR1 ← FX || FEX || VX || OX (if Rc=1)


**fmadd**

Book E	User
--------	------

**fmadd**

**Floating multiply-add [single]**

<b>fmadd</b>	<b>frD,frA,frC,frB</b>	(P=1, Rc=0)
<b>fmadd.</b>	<b>frD,frA,frC,frB</b>	(P=1, Rc=1)
<b>fmadds</b>	<b>frD,frA,frC,frB</b>	(P=0, Rc=0)
<b>fmadds.</b>	<b>frD,frA,frC,frB</b>	(P=0, Rc=1)



$$\text{if } P=1 \text{ then } frD \leftarrow [frA \times_{fp} frC] +_{dp} frB$$

$$\text{else } frD \leftarrow [frA \times_{fp} frC] +_{sp} frB$$

The floating-point operand in **frA** is multiplied by the floating-point operand in **frC**. The floating-point operand in **frB** is added to this intermediate result.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into **frD**.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fmadd[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSAN VXISI VXIMZ  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**fmr**

Book E	User
--------	------

**fmr**

**Floating move register**

**fmr** frD,frB (Rc=0)  
**fmr.** frD,frB (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
1	1	1	1	1	1	frD					///					frB					0	0	0	1	0	0	1	0	0	0	0	Rc

$$frD \leftarrow frB$$

The contents of **frB** are placed into **frD**.

If MSR[FP]=0, an attempt to execute **fmr[.]** causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)

**fmsub**

Book E	User
--------	------

**fmsub**

**Floating multiply-subtract [single]**

**fmsub** frD,frA,frC,frB (P=1, Rc=0)  
**fmsub.** frD,frA,frC,frB (P=1, Rc=1)  
**fmsubs** frD,frA,frC,frB (P=0, Rc=0)  
**fmsubs.** frD,frA,frC,frB (P=0, Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	1	P	1	1	frD					frA					frB					frC					1	1	1	0	0	Rc

$$\begin{aligned} \text{if } P=1 \text{ then } frD &\leftarrow [frA \times_{fp} frC] -_{dp} frB \\ \text{else } frD &\leftarrow [frA \times_{fp} frC] -_{sp} frB \end{aligned}$$

The floating-point operand in **frA** is multiplied by the floating-point operand in **frC**. The floating-point operand in **frB** is subtracted from this intermediate result.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into **frD**.


FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fmsub[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN VXISI VXIMZ  
 CR1 ← FX || FEX || VX || OX (if Rc=1)

<b>fmul</b>	Book E	User	<b>fmul</b>
<b>Floating multiply [single]</b>			
<b>fmul</b>	<b>frD,frA,frC</b>		(P=1, Rc=0)
<b>fmul.</b>	<b>frD,frA,frC</b>		(P=1, Rc=1)
<b>fmuls</b>	<b>frD,frA,frC</b>		(P=0, Rc=0)
<b>fmuls.</b>	<b>frD,frA,frC</b>		(P=0, Rc=1)



if P=1 then frD ← frA ×<sub>dp</sub> frC  
 else frD ← frA ×<sub>sp</sub> frC

The floating-point operand in **frA** is multiplied by the floating-point operand in **frC**.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into **frD**.

Floating-point multiplication is based on exponent addition and multiplication of the significands.


FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fmul[s].[]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN VXIMZ  
 CR1 ← FX || FEX || VX || OX (if Rc=1)

<b>fnabs</b>	Book E	User	<b>fnabs</b>
<b>Floating negative absolute value</b>			
<b>fnabs</b>	<b>frD,frB</b>		(Rc=0)
<b>fnabs.</b>	<b>frD,frB</b>		(Rc=1)



frD ← 0b1||frB<sub>1:63</sub>

The contents of **frB** with bit 0 set are placed into **frD**.

If MSR[FP]=0, an attempt to execute **fnabs.[]** causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)


**fneg**

Book E	User
--------	------

**fneg**

**Floating negate**

**fneg**                                    **frD,frB**                                    (Rc=0)  
**fneg.**                                        **frD,frB**                                    (Rc=1)



$$frD \leftarrow \neg frB_0 || frB_{1:63}$$

The contents of **frB** with bit 0 inverted are placed into **frD**.

If MSR[FP]=0, an attempt to execute **fneg[.]** causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)


**fnmadd**

Book E	User
--------	------

**fnmadd**

**Floating negative multiply-add [single]**

**fnmadd**                                    **frD,frA,frC,frB**                                    (P=1, Rc=0)  
**fnmadd.**                                        **frD,frA,frC,frB**                                    (P=1, Rc=1)  
**fnmadds**                                        **frD,frA,frC,frB**                                    (P=0, Rc=0)  
**fnmadds.**                                        **frD,frA,frC,frB**                                    (P=0, Rc=1)



$$\begin{aligned} \text{if } P=1 \text{ then } frD &\leftarrow -([frA \times_{fp} frC] +_{dp} frB) \\ \text{else } frD &\leftarrow -([frA \times_{fp} frC] +_{sp} frB) \end{aligned}$$

The floating-point operand in **frA** is multiplied by the floating-point operand in **frC**. The floating-point operand in **frB** is added to this intermediate result.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], then negated and placed into **frD**.

This instruction produces the same result as would be obtained by using the Floating Multiply-Add instruction and then negating the result, with the following exceptions.

- QNaNs propagate with no effect on their sign bit.
- QNaNs that are generated as the result of a disabled invalid operation exception have a sign bit of 0.
- SNaNs that are converted to QNaNs as the result of a disabled invalid operation exception retain the sign bit of the SNaN.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

An attempt to execute **fnmadd[s][.]** causes a floating-point unavailable interrupt.





Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN VXISI VXIMZ  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**fnmsub**

Book E	User
--------	------

**fnmsub**

**Floating negative multiply-subtract [single]**

<b>fnmsub</b>	<b>frD,frA,frC,frB</b>	(P=1, Rc=0)
<b>fnmsub.</b>	<b>frD,frA,frC,frB</b>	(P=1, Rc=1)
<b>fnmsubs</b>	<b>frD,frA,frC,frB</b>	(P=0, Rc=0)
<b>fnmsubs.</b>	<b>frD,frA,frC,frB</b>	(P=0, Rc=1)

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

1	1	1	P	1	1	frD	frA	frB	frC	1	1	1	1	0	Rc
---	---	---	---	---	---	-----	-----	-----	-----	---	---	---	---	---	----

if P=1 then frD ← -([frA ×<sub>fp</sub> frC] :<sub>dp</sub> frB)  
else frD ← -([frA ×<sub>fp</sub> frC] :<sub>sp</sub> frB)

The floating-point operand in **frA** is multiplied by the floating-point operand in **frC**. The floating-point operand in **frB** is subtracted from this intermediate result.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], then negated and placed into **frD**.

This instruction produces the same result as would be obtained by using the Floating Multiply-Subtract instruction and then negating the result, with the following exceptions.

- QNaNs propagate with no effect on their sign bit.
- QNaNs that are generated as the result of a disabled invalid operation exception have a sign bit of 0.
- SNaNs that are converted to QNaNs as the result of a disabled invalid operation exception retain the sign bit of the SNaN.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

An attempt to execute **fnmsub[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN VXISI VXIMZ  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**fres**

Book E	User
--------	------

**fres**

**Floating reciprocal estimate single**

<b>fres</b>	<b>frD,frB</b>	(Rc=0)
<b>fres.</b>	<b>frD,frB</b>	(Rc=1)

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

1	1	1	0	1	1	frD	///	frB	///	1	1	0	0	0	Rc
---	---	---	---	---	---	-----	-----	-----	-----	---	---	---	---	---	----

frD ← FPReciprocalEstimate( frB )



A single-precision estimate of the reciprocal of the floating-point operand in **frB** is placed into **frD**. The estimate placed into **frD** is correct to a precision of one part in 256 of the reciprocal of (**frB**), that is,

$$\left| \frac{\text{estimate} - \frac{1}{x}}{\frac{1}{x}} \right| \leq \frac{1}{256}$$

In this example,  $x$  is the initial value in **frB**. Note that the value placed into **frD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the operand is summarized in [Table 204](#).

**Table 204. Operations with special values**

Operand	Result	Exception
$-\infty$	$-0$	None
$-0$	$-\infty$ (No result if FPSCR[ZE] = 1)	ZX
$+0$	$+\infty$ (No result if FPSCR[ZE] = 1)	ZX
$+\infty$	$+0$	None
SNaN	QNaN (No result if FPSCR[VE] = 1.)	VXSNAN
QNaN	QNaN	None

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1 and zero divide exceptions when FPSCR[ZE]=1.

If MSR[FP]=0, an attempt to execute **frs[.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR (undefined) FI (undefined)  
FX OX UX ZX VXSNAN  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**frsp** frsp

Book E	User
--------	------

**Floating round to single-precision**

**frsp** (Rc=0)  
**frsp.** (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
1	1	1	1	1	1					<b>frD</b>											<b>frB</b>	0	0	0	0	0	0	0	1	1	0	0	Rc

If  $\text{frB}[1:11] < 897$  and  $\text{frB}_{1:63} > 0$  then Do  
 If FPSCR[UE] = 0 then goto Disabled Exponent Underflow  
 If FPSCR[UE] = 1 then goto Enabled Exponent Underflow



```

If frB[1:11] > 1150 and frB[1:11] < 2047 then Do
    If FPSCR[OE] = 0 then goto Disabled Exponent Overflow
    If FPSCR[OE] = 1 then goto Enabled Exponent Overflow
If frB[1:11] > 896 and frB[1:11] < 1151 then goto Normal Operand
If frB1:63 = 0 then goto Zero Operand
If frB[1:11] = 2047 then Do
    If frB[12:63] = 0 then goto Infinity Operand
    If frB12 = 1 then goto QNaN Operand
    If frB12 = 0 and frB[13:63] > 0 then goto SNaN Operand
Disabled Exponent Underflow:
    sign ← frB0
    If frB[1:11] = 0 then Do
        exp ← :1022
        frac0:52 ← 0b0 || frB[12:63]
    If frB[1:11] > 0 then Do
        exp ← frB[1:11] : 1023
        fr ← 0b1 || frB[12:63]
Denormalize operand:
    G || R || X ← 0b000
    Do while exp < :126
        exp ← exp + 1
        frac0:52 || G || R || X ← 0b0 || frac0:52 || G || (R |
X)
FPSCR[UX] ← (frac24:52 || G || R || X) > 0
Round Single(sign,exp,frac0:52,G,R,X)
FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
If frac0:52 = 0 then Do
    frD0 ← sign
    frD1:63 ← 0
    If sign = 0 then FPSCR[FPRF] ← '+zero'
    If sign = 1 then FPSCR[FPRF] ← ':zero'
If frac0:52 > 0 then Do
    If frac0 = 1 then Do
        If sign = 0 then FPSCR[FPRF] ← '+normal
number'
        If sign = 1 then FPSCR[FPRF] ← ':normal
number'
    If frac0 = 0 then Do
        If sign = 0 then FPSCR[FPRF] ←
'+denormalized number'
        If sign = 1 then FPSCR[FPRF] ←
':denormalized number'
Normalize operand:
    Do while frac0 = 0
        exp ← exp:1
        frac0:52 ← frac1:52 || 0b0
    frD0 ← sign
    frD[1-11] ← exp + 1023
    frD[12-63] ← frac1:52
Done
Enabled exponent underflow:
    FPSCR[UX] ← 1

```

```

sign ← frB0
If frB[1:11] = 0 then Do
    exp ← :1022
    frac0:52 ← 0b0 || frB[12:63]
If frB[1:11] > 0 then Do
    exp ← frB[1:11] : 1023
    frac0:52 ← 0b1 || frB[12:63]
Normalize operand:
    Do while frac0 = 0
        exp ← exp : 1
        frac0:52 ← frac1:52 || 0b0
Round Single(sign,exp,frac0:52,0,0,0)
FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
exp ← exp + 192
frD0 ← sign
frD[1-11] ← exp + 1023
frD[12-63] ← frac1:52
If sign = 0 then FPSCR[FPRF] ← '+normal number'
If sign = 1 then FPSCR[FPRF] ← ':normal number'
Done
Disabled exponent overflow
FPSCR[OX] ← 1
If FPSCR[RN] = 0b00 then Do /* Round to Nearest
*/
    If frB0 = 0 then frD ←
0x7FF0_0000_0000_0000
    If frB0 = 1 then frD ←
0xFFFF0_0000_0000_0000
    If frB0 = 0 then FPSCR[FPRF] ← '+infinity'
    If frB0 = 1 then FPSCR[FPRF] ← ':infinity'
If FPSCR[RN] = 0b01 then Do /* Round toward
Zero */
    If frB0 = 0 then frD ←
0x47EF_FFFF_E000_0000
    If frB0 = 1 then frD ←
0xC7EF_FFFF_E000_0000
    If frB0 = 0 then FPSCR[FPRF] ← '+normal
number'
    If frB0 = 1 then FPSCR[FPRF] ← ':normal
number'
If FPSCR[RN] = 0b10 then Do /* Round toward
+Infinity */
    If frB0 = 0 then frD ←
0x7FF0_0000_0000_0000
    If frB0 = 1 then frD ←
0xC7EF_FFFF_E000_0000
    If frB0 = 0 then FPSCR[FPRF] ← '+infinity'
    If frB0 = 1 then FPSCR[FPRF] ← ':normal
number'
If FPSCR[RN] = 0b11 then Do /* Round toward
:Infinity */
    If frB0 = 0 then frD ←

```

```

0x47EF_FFFF_E000_0000
0xFFFF0_0000_0000_0000
number'
    If frB0 = 1 then frD ←
    If frB0 = 0 then FPSCR[FPRF] ← '+normal
    If frB0 = 1 then FPSCR[FPRF] ← ':infinity'
    FPSCR[FR] ← undefined
    FPSCR[FI] ← 1
    FPSCR[XX] ← 1
    Done
Enabled Exponent Overflow:
    sign ← frB0
    exp ← frB[1:11] : 1023
    frac0:52 ← 0b1 || frB[12:63]
    Round Single(sign,exp,frac0:52,0,0,0)
    FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
Enabled Overflow:
    FPSCR[OX] ← 1
    exp ← exp : 192
    frD0 ← sign
    frD[1-11] ← exp + 1023
    frD[12-63] ← frac1:52
    If sign = 0 then FPSCR[FPRF] ← '+normal number'
    If sign = 1 then FPSCR[FPRF] ← ':normal number'
    Done
Zero Operand:
    frD ← frB
    If frB0 = 0 then FPSCR[FPRF] ← '+zero'
    If frB0 = 1 then FPSCR[FPRF] ← ':zero'
    FPSCR[FR,FI] ← 0b00
    Done
Infinity Operand:
    frD ← frB
    If frB0 = 0 then FPSCR[FPRF] ← '+infinity'
    If frB0 = 1 then FPSCR[FPRF] ← ':infinity'
    FPSCR[FR,FI] ← 0b00
    Done
QNaN Operand:
    frD ← frB0:34 || 290
    FPSCR[FPRF] ← 'QNaN'
    FPSCR[FR,FI] ← 0b00
    Done
SNaN Operand:
    FPSCR[VXSNAN] ← 1
    If FPSCR[VE] = 0 then Do
        frD[0:11] ← frB[0:11]
        frD12 ← 1
        frD[13:63] ← frB[13:34] || 290
        FPSCR[FPRF] ← 'QNaN'
    FPSCR[FR,FI] ← 0b00
    Done
Normal Operand:

```

```

sign ← frB0
exp ← frB[1:11] : 1023
frac0:52 ← 0b1 || frB[12:63]
Round Single(sign,exp,frac0:52,0,0,0)
FPSCR[XX] ← FPSCR[XX] | FPSCR[FI]
If exp > 127 and FPSCR[OE] = 0 then go to Disabled
Exponent Overflow
If exp > 127 and FPSCR[OE] = 1 then go to Enabled
Overflow
frD0 ← sign
frD[1-11] ← exp + 1023
frD[12-63] ← frac1:52
If sign = 0 then FPSCR[FPRF] ← '+normal number'
If sign = 1 then FPSCR[FPRF] ← ':normal number'
Done
Round Single(sign,exp,frac0:52,G,R,X):
inc ← 0
lsb ← frac23
gbit ← frac24
rbit ← frac25
xbit ← (frac26:52 || G || R || X) ≠ 0
If FPSCR[RN] = 0b00 then Do /* comparison ignores u
bits */
If sign || lsb || gbit || rbit || xbit = 0bu11uu then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0bu011u then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0bu01u1 then
inc ← 1
If FPSCR[RN] = 0b10 then Do /* comparison ignores u
bits */
If sign || lsb || gbit || rbit || xbit = 0b0u1uu then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0b0uu1u then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0b0uuu1 then
inc ← 1
If FPSCR[RN] = 0b11 then Do /* comparison ignores u
bits */
If sign || lsb || gbit || rbit || xbit = 0b1u1uu then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0b1uu1u then
inc ← 1
If sign || lsb || gbit || rbit || xbit = 0b1uuu1 then
inc ← 1
frac0:23 ← frac0:23 + inc
If carry_out = 1 then Do
frac0:23 ← 0b1 || frac0:22
exp ← exp + 1
frac24:52 ← 290
FPSCR[FR] ← inc

```

FPSCR[FI] ← gbit | rbit | xbit  
Return

The floating-point operand in **frB** is rounded to single-precision, using the rounding mode specified by FPSCR[RN], and placed into **frD**.


FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **frsp[.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX OX UX XX VXSNaN  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**frsqrte** Book E User **frsqrte**  
**Floating reciprocal square root estimate**  
**frsqrte** frD,frB (Rc=0)  
**frsqrte.** frD,frB (Rc=1)



frD ← FPReciprocalSquareRootEstimate( frB )

A double-precision estimate of the reciprocal of the square root of the floating-point operand in **frB** is placed into **frD**. The estimate is correct to a precision of one part in 32 of the reciprocal of the square root of (**frB**), that is,

Here, x is the initial value in **frB**. Note that the value placed into **frD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the operand is summarized in [Table 205](#)

**Table 205. Operations with special values**

Operand	Result	Exception
$-\infty$	QNaN (No result if FPSCR[VE] = 1.)	VXSQRT
< 0	QNaN (No result if FPSCR[VE] = 1.)	VXSQRT
-0	$-\infty$ (No result if FPSCR[ZE] = 1.)	ZX
+0	$+\infty$ (No result if FPSCR[ZE] = 1.)	ZX
$+\infty$	+0	None
SNaN	QNaN (No result if FPSCR[VE] = 1.)	VXSNaN
QNaN	QNaN	None

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1 and zero divide exceptions when FPSCR[ZE]=1.

If MSR[FP]=0, attempting to execute **frsqrte[.]** causes a floating-point unavailable interrupt.


Other registers altered:

- FPRF FR (undefined) FI (undefined)  
FX ZX VXSNaN VXSQRT  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**fsel** Book E User **fsel**

**Floating select**

**fsel** frD,frA,frC,frB (Rc=0)  
**fsel.** frD,frA,frC,frB (Rc=1)



if frA ≥ 0.0 then frD ← frC  
else frD ← frB

The floating-point operand in **frA** is compared to the value zero. If the operand is greater than or equal to zero, **frD** is set to the contents of **frC**. If the operand is less than zero or is a NaN, **frD** is set to the contents of **frB**. The comparison ignores the sign of zero (that is, +0 and -0 are regarded as equal).

If MSR[FP]=0, an attempt to execute **fsel**[.] causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)

*Note:* Programming: Examples of uses of this instruction can be found in the appendix

---


**Warning: Care must be taken in using fsel if IEEE compatibility is required, or if the values being tested can be NaNs or infinities**

---

**fsqrt** Book E User **fsqrt**

**Floating square root [single]**

**fsqrt** frD,frB (P=1, Rc=0)  
**fsqrt.** frD,frB (P=1, Rc=1)  
**fsqrts** frD,frB (P=0, Rc=0)  
**fsqrts.** frD,frB (P=0, Rc=1)



if P=1 then frD ← FPSquareRootDouble( frB )  
else frD ← FPSquareRootSingle( frB )

The square root of the floating-point operand in **frB** is placed into **frD**.





If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into frD.

Operation with various special values of the operand is summarized in [Table 206](#)

**Table 206. Operations with special values**

Operand	Result	Exception
$-\infty$	QNaN (No result if FPSCR[VE] = 1)	VXSQRT
< 0	QNaN (No result if FPSCR[VE] = 1)	VXSQRT
-0	-0	None
$+\infty$	$+\infty$	None
SNaN	QNaN(No result if FPSCR[VE] = 1)	VXSNAN
QNaN	QNaN	None

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fsqrt[s][.]** causes a floating-point unavailable interrupt.

Other registers altered:

- FPRF FR FI FX XX VXSNAN VXSQRT  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**fsub**

Book E	User
--------	------

**fsub**

**Floating subtract [single]**

<b>fsub</b>	<b>frD,frA,frB</b>	(P=1, Rc=0)
<b>fsub.</b>	<b>frD,frA,frB</b>	(P=1, Rc=1)
<b>fsubs</b>	<b>frD,frA,frB</b>	(P=0, Rc=0)
<b>fsubs.</b>	<b>frD,frA,frB</b>	(P=0, Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	1	P	1	1	frD			frA			frB			///			1	0	1	0	0	Rc								

if P=1 then frD ← frA <sub>-dp</sub> frB  
else frD ← frA <sub>-sp</sub> frB

The floating-point operand in frB is subtracted from the floating-point operand in frA.

If the msb of the resultant significand is not 1, the result is normalized. The result is rounded to the target precision under control of the floating-point rounding control field, FPSCR[RN], and placed into frD.

The execution of the Floating Subtract instruction is identical to that of Floating Add, except that the contents of frB participate in the operation with the sign bit (bit 0) inverted.

FPSCR[FPRF] is set to the class and sign of the result, except for invalid operation exceptions when FPSCR[VE]=1.

If MSR[FP]=0, an attempt to execute **fsub[s][.]** causes a floating-point unavailable interrupt.



Other registers altered:

- FPRF FR FI FX OX UX XX VXSNNAN VXISI  
CR1 ← FX || FEX || VX || OX (if Rc=1)

**icbi** Book E User **icbi**

**Instruction cache block invalidate**

**icbi** rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	///					rA	rB	1 1 1 1 0 1 0 1 1 0					/													

if rA=0 then a ← <sup>64</sup>0 else a ← rA  
EA ← <sup>32</sup>0 || (a + rB)<sub>32:63</sub>  
InvalidateInstructionCacheBlock( EA )

EA calculation:      Addressing ModeEA for rA=0EA for rA≠0  
<sup>32</sup>0 || rB<sub>32:63</sub> <sup>32</sup>0 || (rA+rB)<sub>32:63</sub>

If the block containing the byte addressed by EA is in memory that is memory-coherence required and a block containing the byte addressed by EA is in the instruction cache of any processors, the block is invalidated in those instruction caches, so that subsequent references cause the block to be fetched from main memory.

If the block containing the byte addressed by EA is in memory that is not memory-coherence required and a block containing the byte addressed by EA is in the instruction cache of this processor, the block is invalidated in that instruction cache, so that subsequent references cause the block to be fetched from main memory.

The function of this instruction is independent of whether the block containing the byte addressed by EA is in memory that is write-through required or caching-inhibited.

This instruction is treated as a load.

**icbi** may cause a cache-locking exception on some implementations. See the implementation documentation.

On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

Other registers altered: None

**icblc** Cache locking APU User **icblc**

**Instruction cache block lock clear**

**icblc** CT,rA,rB Form: X

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CT					rA	rB	0 0 1 1 1 0 0 1 1 0					/													

if rA = 0 then a ← <sup>64</sup>0 else a ← GPR(rA)  
if Mode32 then EA ← <sup>32</sup>0 || (a + GPR(rB))<sub>32:63</sub>  
if Mode64 then EA ← a + GPR(rB)  
InstructionCacheBlockClearLock(CT, EA)

EA calculation:      EA for rA=0EA for rA≠0



$$^{32}0 \ || \ GPR(rB)_{32:63} \ ^{32}0 \ || \ (GPR(rA) + GPR(rB))_{32:63}$$

The instruction cache specified by CT has the cache line corresponding to EA unlocked allowing the line to participate in the normal replacement policy.

Cache lock clear instructions remove locks previously set by cache lock set instructions.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

The **icbtlc** instruction requires read (R) or execute (X) permissions with respect to translation and memory protection and can cause DSI and DTLB error interrupts accordingly.

An unable-to-unlock condition is said to occur any of the following conditions exist:

- The target address is marked cache-inhibited, or the storage attributes of the address uses a coherency protocol that does not support locking.
- The target cache is disabled or not present.
- The CT field of the instructions contains a value not supported by the implementation.
- The target address is not in the cache or is present in the cache but is not locked.

If an unable-to-unlock condition occurs, no cache operation is performed.

**EIS specifics**

Setting L1CSR1[ICLFI] allows system software to clear all L1 instruction cache locking bits without knowing the addresses of the lines locked.

**icbt** Book E User **icbt**

**Instruction cache block touch**

**icbt** CT,rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	1	1	1	1	1	CT						rA						rB						0	0	0	0	0	1	0	1	1	0	/

if rA=0 then a ← <sup>64</sup>0 else a ← rA  
 EA ← <sup>32</sup>0 || (a + rB)<sub>32:63</sub>  
 PrefetchInstructionCacheBlock( CT, EA )

EA calculation:                      Addressing ModeEA for rA=0EA for rA≠0  
<sup>32</sup>0 || rB<sub>32:63</sub> <sup>32</sup>0 || (rA+rB)<sub>32:63</sub>

This instruction is a hint that performance would likely be improved if the block containing the byte addressed by EA is fetched into the instruction cache, because the program will probably soon execute code from the addressed location.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

Implementations should perform no operation when CT specifies a value not supported by the implementation.

The hint is ignored if the block is caching-inhibited.



This instruction treated as a load (see the discussion of cache and MMU operation in the user’s manual), except that an interrupt is not taken for a translation or protection violation.

Other registers altered: None

**icbtl** Cache locking APU User **icbtl**

**Instruction cache block touch and lock set**

**icbtl** CT,rA,rB Form: X

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1					CT												0	1	1	1	1	0	0	1	1	0	/

if rA = 0 then a ← 640 else a ← GPR(rA)  
 if Mode32 then EA ← 320 || (a + GPR(rB))32:63  
 if Mode64 then EA ← a + GPR(rB)  
 PrefetchInstructionCacheBlockLockSet(CT, EA)  
 EA calculation: EA for rA=0EA for rA≠0  
 320 || GPR(rB)32:63 320 || (GPR(rA) + GPR(rB))32:63

The instruction cache specified by CT has the line corresponding to EA loaded and locked. If the line exists in the cache, it is locked without refetching from memory.

Cache touch and lock set instructions allow software to lock lines into the cache to shorten latency for critical cache accesses and more deterministic behavior. Lines locked in the cache do not participate in the normal replacement policy when a line must be victimized for replacement.

[Section 4.3.1.18: User-level cache instructions](#), lists supported CT values. An implementation may use other CT values to enable software to target specific, implementation-dependent portions of its cache hierarchy or structure.

The **icbtl** requires read (R) or execute (X) permissions for translation and memory protection and can cause DSI and DTLB error interrupts accordingly.

For unable-to-lock conditions, described in [Section 9.1.1.4: Unable-to-lock conditions](#), no cache operation is performed and LICSR0[ICUL] is set.

An overlocking condition is said to exist is all the available ways for a given cache index are already locked. If an overlocking condition occurs for a **icbtl** instruction and if the lock was targeted for the primary cache or secondary cache (CT = 0 or CT = 2), the requested line is not locked into the cache. When an overlock condition occurs, L1CSR1[ICLO] (L2CSR[L2CLO] for CT = 2) is set. If L1CSR1[ICLOA] is set (or L2CSR[L2CLOA] for CT = 2), the requested line is locked into the cache and implementation dependent line currently locked in the cache is evicted.

Results of overlocking and unable-to-lock conditions for caches other than the primary and secondary cache are defined as part of the architecture for the cache hierarchy designated by CT.

If a unified primary cache is implemented and L1CSR1 is not implemented, L1CSR0[DCUL] and L1CSR0[DCLO] are updated instead of the corresponding L1CSR1 bits.

Other registers altered:

- L1CSR1[ICUL] if unable to lock occurs
- L1CSR1[ICLO] (L2CSR[L2CLO]) if lock overflow occurs

**\_illegal**

VLE	User
-----	------

**\_illegal**

**Illegal**

**se\_illegal**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRR1 ← MSR

SRR0 ← CIA

NIA ← IVPR<sub>32:47</sub> || IVOR<sub>648:59</sub> || 0b0000

MSR<sub>WE,EE,PR,IS,DS,FP,FE0,FE1</sub> ← 0b0000\_0000

**se\_illegal** is used to request an illegal instruction exception. A program interrupt is generated. The contents of the MSR are copied into SRR1 and the address of the **se\_illegal** instruction is placed into SRR0.

MSR[WE,EE,PR,IS,DS,FP,FE0,FE1] are cleared.

The interrupt causes the next instruction to be fetched from address IVPR[32–47]||IVOR6[48–59]||0b0000

This instruction is context synchronizing.

Special Registers Altered: SRR0 SRR1 MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]

**isel**

Integer Select APU	User
--------------------	------

**isel**

**Integer Select**

**isel** rD, rA, rB, crb

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1																						0	1	1	1	1	0

if (rA = 0) then a ← 640 else a ← GPR(rA)  
 c ← CR<sub>crb + 32</sub>  
 if c then rD ← a  
 else rD ← GPR(rB)

If CR[crb + 32] is set, the contents of rA|0 are copied into rD. If CR[crb + 32] is clear, the contents of rB are copied into rD.

**isync**

Book E	User
--------	------

**isync**

**Instruction synchronize**

**isync**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	0	0	1	1																										

**isync** provides an ordering function for the effects of all instructions executed by the processor executing the **isync** instruction. Executing an **isync** ensures that all instructions

preceding the **isync** have completed before **isync** completes, and that no subsequent instructions are initiated until after **isync** completes. It also causes any prefetched instructions to be discarded, with the effect that subsequent instructions are fetched and executed in the context established by the instructions preceding **isync**.

**isync** may complete before memory accesses associated with instructions preceding **isync** have been performed.

**isync** is context synchronizing. See [Section 4.2.3.6: Context synchronization](#).

Other registers altered: None


**\_isync**

VLE	User
-----	------

**\_isync**

**Instruction Synchronize**

**se\_isync**



The **se\_isync** instruction provides an ordering function for the effects of all instructions executed by the processor executing the **se\_isync** instruction. Executing an **se\_isync** instruction ensures that all instructions preceding the **se\_isync** instruction have completed before the **se\_isync** instruction completes, and that no subsequent instructions are initiated until after the **se\_isync** instruction completes. It also causes any prefetched instructions to be discarded, with the effect that subsequent instructions are fetched and executed in the context established by the instructions preceding the **se\_isync** instruction.

The **se\_isync** instruction may complete before memory accesses associated with instructions preceding the **se\_isync** instruction have been performed.

This instruction is context synchronizing (see Book E). It has identical semantics to Book E **isync**, just a different encoding.

Special Registers Altered: None


**lbz**

Book E	User
--------	------


**lbz**

**Load byte and zero [with update] [indexed]**

**lbz**                    rD,D(rA)                    (D-mode, U=0)  
**lbzu**                    rD,D(rA)                    (D-mode, U=1)



**lbzx**                    rD,rA,rB                    (X-mode, U=0)  
**lbzux**                    rD,rA,rB                    (X-mode, U=1)



if rA=0 then a ← <sup>64</sup>0 else a ← rA  
 if D-mode then EA ← <sup>32</sup>0 || (a + EXTS(D))<sub>32:63</sub>  
 if X-mode then EA ← <sup>32</sup>0 || (a + rB)<sub>32:63</sub>



$rD \leftarrow {}^{56}0 \parallel \text{MEM}(\text{EA}, 1)$   
 if  $U=1$  then  $rA \leftarrow \text{EA}$

The EA is calculated as follows:

- For **lbz** and **lbzu**, EA is bits 32–63 of the sum of the contents of  $rA$ , or 64 zeros if  $rA=0$ , and the sign-extended value of the D field.
- For **lbzx** and **lbzux**, EA is bits 32–63 of the sum of the contents of  $rA$ , or 64 zeros if  $rA=0$ , and the contents of  $rB$ .

The byte in memory addressed by EA is loaded into  $rD[56–63]$ ;  $rD[0–55]$  are cleared.


If  $U=1$  (with update), EA is placed into  $rA$ . If  $U=1$  (with update), and  $rA=0$  or  $rA=rD$ , the instruction form is invalid.

Other registers altered: None


**\_lbzx** VLE User **\_lbzx**

**Load Byte and Zero [with Update] [Indexed]**


**e\_lbz**  $rD, D(rA)$  (D-mode)



**se\_lbz**  $rZ, SD4(rX)$  (SD4-mode)



**e\_lbzu**  $rD, D8(rA)$  (D8-mode)



if  $(RA=0 \ \& \ !se\_lbz)$  then  $a \leftarrow {}^{32}0$  else  $a \leftarrow \text{GPR}(RA \text{ or } RX)$

if D-mode then  $\text{EA} \leftarrow (a + \text{EXTS}(D))_{32:63}$

if D8-mode then  $\text{EA} \leftarrow (a + \text{EXTS}(D8))_{32:63}$

if SD4-mode then  $\text{EA} \leftarrow (a + ({}^{28}0 \parallel \text{SD4}))_{32:63}$

$\text{GPR}(RD \text{ or } RZ) \leftarrow {}^{24}0 \parallel \text{MEM}(\text{EA}, 1)$

if  $e\_lbzu$  then  $\text{GPR}(RA) \leftarrow \text{EA}$

Let the EA be calculated as follows:

- For **e\_lbz** and **e\_lbzu**, let EA be the sum of the contents of  $\text{GPR}(rA)$ , or 32 0s if  $rA = 0$ , and the sign-extended value of the D or D8 instruction field.
- For **se\_lbz**, let EA be the sum of the contents of  $\text{GPR}(rX)$  and the zero-extended value of the SD4 instruction field.

The byte in memory addressed by EA is loaded into bits 56–63 of  $\text{GPR}(rD \text{ or } rZ)$ . Bits 32–55 of  $\text{GPR}(rD \text{ or } rZ)$  are cleared.

If **e\_lbzu**, EA is placed into  $\text{GPR}(rA)$ .

If **e\_lbzu** and  $rA = 0$  or  $rA = rD$ , the instruction form is invalid.

Special Registers Altered: None


**lfd**

Book E	User
--------	------


**lfd**

**Load floating-point double**

**lfd** frD,D(rA) (D-mode, U=0)  
**lfdU** frD,D(rA) (D-mode, U=1)



**lfdx** frD,rA,rB (X-mode, U=0)  
**lfdxU** frD,rA,rB (X-mode, U=1)



if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))32:63  
 if X-mode then EA ← 320 || (a + rB)32:63  
 frD ← MEM(EA,8)  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **lfd** and **lfdU**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **lfdx** and **lfdxU**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The double word addressed by EA is placed into frD.

If U=1 (with update), EA is placed into register rA.

If U=1 (with update) and rA=0, the instruction form is invalid.

If MSR[FP]=0, an attempt to execute **lfd[u][x]** causes a floating-point unavailable interrupt.

Other registers altered: None


**lfs**

Book E	User
--------	------

**lfs**

**Load floating-point single**

**lfs** frD,D(rA) (D-mode, U=0)  
**lfsU** frD,D(rA) (D-mode, U=1)





**lfsx** **frD,rA,rB** (X-mode, U=0)  
**lfsux** **frD,rA,rB** (X-mode, U=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	frD					rA					rB					1	0	0	0	U	1	0	1	1	1	/

if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))<sub>32:63</sub>  
 if X-mode then EA ← 320 || (a + rB)<sub>32:63</sub>  
 frD ← DOUBLE(MEM(EA,4))  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **lfs** and **lfsu**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **lfsx** and **lfsux**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The word addressed by EA is interpreted as a single-precision operand, converted to floating-point double format, and placed into frD.

If U=1 (with update), EA is placed into register rA.

If U=1 (with update) and rA=0, the instruction form is invalid.

If MSR[FP]=0, an attempt to execute **lfs[u][x]** causes a floating-point unavailable interrupt.

Other registers altered: None

**lha** Book E User **lha**

**Load half word algebraic [with update] [indexed]**

**lha** **rD,D(rA)** (D-mode, U=0)  
**lhau** **rD,D(rA)** (D-mode, U=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	0	1	0	1	U	rD					rA					D															

**lhax** **rD,rA,rB** (X-mode, U=0)  
**lhaux** **rD,rA,rB** (X-mode, U=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					rA					rB					0	1	0	1	U	1	0	1	1	1	/

if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))<sub>32:63</sub>  
 if X-mode then EA ← 320 || (a + rB)<sub>32:63</sub>  
 rD ← 320 || EXTS(MEM(EA,2))<sub>32:63</sub>  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **lha** and **lhau**, EA is bits 32–63 of the sum of the contents of **rA**, or 64 zeros if **rA=0**, and the sign-extended value of the D field.
- For **lhax** and **lhaux**, EA is bits 32–63 of the sum of the contents of **rA**, or 64 zeros if **rA=0**, and the contents of **rB**.

The half word addressed by EA is loaded into **rD[48–63]**. **rD[32–47]** are filled with a copy of bit 0 of the loaded half word. Bits **rD[0–31]** are cleared.

If **U=1** (with update), EA is placed into **rA**.


If **U=1** (with update), and **rA=0** or **rA=rD**, the instruction form is invalid.

Other registers altered: None


**\_lhax** VLE User **\_lhax**

**Load Halfword Algebraic [with Update] [Indexed]**

**e\_lha** **rD,D(rA)** (D-mode)



**e\_lhau** **rD,D8(rA)** (D8-mode)



if RA=0 then a ← 320 else a ← GPR(RA)  
 if D-mode then EA ← (a + EXTS(D))32:63  
 if D8-mode then EA ← (a + EXTS(D8))32:63  
 GPR(RD) ← EXTS(MEM(EA,2))32:63  
 if e\_lhau then GPR(RA) ← EA

Let the EA be calculated as follows:

- For **e\_lha** and **e\_lhau**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.

The half word in memory addressed by EA is loaded into bits 48–63 of GPR(rD). Bits 32–47 of GPR(rD) are filled with a copy of bit 0 of the loaded half word.

If **e\_lhau**, EA is placed into GPR(rA).


If **e\_lhau** and rA = 0 or rA = rD, the instruction form is invalid.

Special Registers Altered: None

**lhbrx** Book E User **lhbrx**

**Load half word byte-reverse indexed**

**lhbrx** **rD,rA,rB**



if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)32:63  
 data0:15 ← MEM(EA,2)  
 rD ← 480 || data8:15 || data0:7

The EA is calculated as follows:

- For **lhbrx**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

Bits 0–7 of the half word addressed by EA are loaded into rD[56–63]. Bits 8–15 of the half word addressed by EA are loaded into rD[48–55]; rD[0–47] are cleared.


Other registers altered: None




Programming notes:

- When EA references big-endian memory, these instructions have the effect of loading data in little-endian byte order. Likewise, when EA references little-endian memory, these instructions have the effect of loading data in big-endian byte order.
- In some implementations, the Load Half Word Byte-Reverse Indexed instructions may have greater latency than other load instructions.

<b>lhz</b>	Book E	User	<b>lhz</b>
<b>Load half word and zero [with update] [indexed]</b>			
<b>lhz</b>	rD,D(rA)		(D-mode, U=0)
<b>lhzu</b>	rD,D(rA)		(D-mode, U=1)



<b>lhzx</b>			<b>lhzx</b>
<b>lhzx</b>	rD,rA,rB		(X-mode, U=0)
<b>lhzux</b>	rD,rA,rB		(X-mode, U=1)



if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))<sub>32:63</sub>  
 if X-mode then EA ← 320 || (a + rB)<sub>32:63</sub>  
 rD ← 480 || MEM(EA,2)  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **lhz** and **lhzu**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **lhzx** and **lhzux**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.


The half word addressed by EA is loaded into rD[48–63]; rD[0–47] are cleared.

If U=1 (with update), EA is placed into rA.


If U=1 (with update), and rA=0 or rA=rD, the instruction form is invalid.

Other registers altered: None


<b>_lhzx</b>	VLE	User	<b>_lhzx</b>
<b>Load Halfword and Zero [with Update] [Indexed]</b>			
<b>e_lhz</b>	rD,D(rA)		(D-mode)



**se\_lhz** **rZ,SD4(rX)** (SD4-mode)



**e\_lhzu** **rD,D8(rA)** (D8-mode)



if (RA=0 & !se\_lhz) then a ← <sup>32</sup>0 else a ← GPR(RA or RX)  
 if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>  
 if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>  
 if SD4-mode then EA ← (a + (<sup>27</sup>0 || SD4 || 0))<sub>32:63</sub>  
 GPR(RD or RZ) ← <sup>16</sup>0 || MEM(EA,2)  
 if e\_lhzu then GPR(RA) ← EA

Let the EA be calculated as follows:


- For **e\_lhz** and **e\_lhzu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_lhz** let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 1 bit.

The half word in memory addressed by EA is loaded into bits 48–63 of GPR(rD). Bits 32–47 of GPR(rD) are cleared.

If **e\_lhzu**, EA is placed into GPR(rA).


If **e\_lhzu** and rA = 0 or rA = rD, the instruction form is invalid.

**Special Registers Altered: None**



**Load Immediate [Shifted]**

**e\_li** **rD,LI20** (LI20-mode)




LI20 ← LI20<sub>0:3</sub> || LI20<sub>4:8</sub> || LI20<sub>9:19</sub>  
 GPR(RD) ← EXTS(LI20)

For **e\_li**, the sign-extended LI20 field is placed into GPR(rD).

**Special Registers Altered: None**

**e\_lis** **rD,UI**




UI ← UI<sub>0:4</sub> || UI<sub>5:15</sub>  
 GPR(RD) ← UI || <sup>16</sup>0



For **e\_lis**, the UI field is concatenated on the right with 16 0's and placed into GPR(**rD**).

Special Registers Altered: None

**se\_li**



$$\text{GPR}(\text{RX}) \leftarrow {}^{25}0 \parallel \text{UI7}$$

For **se\_li**, the zero-extended UI7 field is placed into GPR(**rX**).

Special Registers Altered: None

**Imw**


Book E	User
--------	------

**Imw**

**Load multiple word**

**Imw**

**rD,D(rA)**



if  $rA=0$  then  $EA \leftarrow {}^{32}0 \parallel \text{EXTS}(D)_{32:63}$   
 else  $EA \leftarrow {}^{32}0 \parallel (rA+\text{EXTS}(D))_{32:63}$   
 $r \leftarrow rD$   
 do while  $r \leq 31$

$\text{GPR}(r) \leftarrow {}^{32}0 \parallel \text{MEM}(EA,4)$   
 $r \leftarrow r + 1$   
 $EA \leftarrow {}^{32}0 \parallel (EA+4)_{32:63}$

The EA is bits 32–63 of the sum of the contents of **rA**, or 64 zeros if **rA**=0, and the sign-extended value of the D instruction field.

Here  $n=(32-rD)$ . n consecutive words starting at EA are loaded into bits 32–63 of registers **rD** through GPR31. Bits 0–31 of these GPRs are cleared.

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined. If **rA** is in the range of registers to be loaded, including the case in which **rA**=0, the instruction form is invalid.

Other registers altered: None

**\_Imw**


VLE	User
-----	------

**\_Imw**

**Load Multiple Word**

**e\_Imw**

**rD,D8(rA)**



if  $RA=0$  then  $EA \leftarrow \text{EXTS}(D8)_{32:63}$   
 else  $EA \leftarrow (\text{GPR}(RA)+\text{EXTS}(D8))_{32:63}$   
 $r \leftarrow RD$

```
do while r ≤ 31
    GPR(r) ← MEM(EA,4)
    r ← r + 1
    EA ← (EA+4)32:63
```



Let the EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D8 instruction field.

Let n = (32-rD). n consecutive words starting at EA are loaded into bits 32–63 of registers GPR(rD) through GPR(31).

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined. If rA is in the range of registers to be loaded, including the case in which rA = 0, the instruction form is invalid.

Special Registers Altered: None

**lswi** Book E User **lswi**  
**Load string word (immediate | indexed)**



```
if rA=0 then a ← 640 else a ← rA
if 'lswi' then EA ← 320 || a32:63
if 'lswx' then EA ← 320 || (a + rB)32:63
if 'lswi' & NB=0 then n ← 32
if 'lswi' & NB≠0 then n ← NB
if 'lswx' then n ← XER57:63
r ← rD : 1
i ← 32
rD ← undefined
do while n > 0
    if i = 32 then
        r ← r + 1 (mod 32)
        GPR(r) ← 0
        GPR(r)i:i+7 ← MEM(EA,1)
        i ← i + 8
    if i = 64 then i ← 32
    EA ← 320 || (EA+1)32:63
    n ← n : 1
```

The EA is calculated as follows:

- For **lswi**, EA is 32 zeros concatenated with rA[32–63], or 32 zeros if rA=0.
- For **lswx**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.



If **lswi**,  $n = NB$  if  $NB \neq 0$ ,  $n = 32$  if  $NB=0$ . If **lswx**,  $n=XER[57-63]$ .  $n$  is the number of bytes to load. Here  $nr=CEIL(n\div 4)$ :  $nr$  is the number of registers to receive data.

If  $n>0$ ,  $n$  consecutive bytes starting at  $EA$  are loaded into registers  $rD$  through  $(rD+nr-1)$ . Data is loaded into the low-order 4 bytes of each GPR; the high-order 4 bytes are cleared.

Bytes are loaded left to right in each GPR. The sequence wraps to GPR0 if required. If the 4 LSBs of GPR( $rD+nr-1$ ) are partially filled, the unfilled LSBs of that GPR are cleared.

If **lswx** and  $n=0$ , the contents of  $rD$  are undefined.

If  $rA$ , or  $rB$  for **lswx**, is in the range of registers to be loaded, including where  $rA=0$ , an illegal instruction type program interrupt is invoked or results are boundedly undefined. If  $rD=rA$ , or  $rD=rB$  for **lswx**, the instruction form is invalid. Other registers altered: None

*Note: Programming: String instructions move data without concern for alignment. They can perform short moves between arbitrary locations or long moves between misaligned memory fields.*


**lwarx**

Book E	User
--------	------

**lwarx**

**Load word and reserve indexed**

**lwarx**  $rD,rA,rB$



```

if rA=0 then a ← 640 else a ← rA
EA ← 320 || (a + rB)32:63
RESERVE ← 1
RESERVE_ADDR ← real_addr(EA)
rD ← 320 || MEM(EA,4)
    
```

$EA$  is bits 32–63 of the sum of the contents of  $rA$  (32 zeros if  $rA=0$ ), and the contents of  $rB$ . The word addressed by  $EA$  is loaded into  $rD[32-63]$ ;  $rD[0-31]$  are cleared.

**lwarx** creates a reservation for use by a **stwcx**. instruction. An address computed from the  $EA$  is associated with the reservation and replaces any previously associated address. See [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#)

If  $EA$  is not a multiple of 4, an alignment interrupt occurs or results are boundedly undefined.

Other registers altered: None

Programming notes:

- **lwarx**, and **stwcx**. permit programmers to write an instruction sequence that appears to perform an atomic update operation on a memory location. This operation depends on a single reservation resource in each processor. At most one reservation exists on any given processor.
- Because **lwarx** instructions have implementation dependencies (such as the granularity at which reservations are managed), they must be used with care. System library programs should use these instructions to implement high-level synchronization functions (such as test and set, compare and swap) needed by application programs. Application programs should use these library programs, rather than use **lwarx** directly. The granularity with which reservations are managed is implementation-dependent.






Therefore the location to be accessed by **lwarx** should be allocated by a system library program. See [Section 4.3.1.16: Atomic update primitives using lwarx and stwxc](#).

**lwbx** Book E User **lwbx**

**Load word byte-reverse indexed**

**lwbx** rD,rA,rB



if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)<sub>32:63</sub>  
 data<sub>0:31</sub> ← MEM(EA,4)  
 rD ← 320 || data<sub>24:31</sub> || data<sub>16:23</sub> || data<sub>8:15</sub> || data<sub>0:7</sub>

The EA is calculated as follows:

- For **lwbx**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

Bits 0–7 of the word addressed by EA are loaded into rD[56–63]. Bits 8–15 of the word addressed by EA are loaded into rD[48–55]. Bits 16–23 of the word addressed by EA are loaded into rD[40–47]. Bits 24–31 of the word addressed by EA are loaded into rD[32–39]. Bits rD[0–31] are cleared.

Other registers altered: None


Programming notes:

- When EA references big-endian memory, these instructions have the effect of loading data in little-endian byte order. Likewise, when EA references little-endian memory, these instructions have the effect of loading data in big-endian byte order.
- In some implementations, the load word byte-reverse instructions may have greater latency than other load instructions.


**lwz** Book E User **lwz**

Load word and zero [with update] [indexed]

**lwz** rD,D(rA) (D-mode, U=0)  
**lwzu** rD,D(rA) (D-mode, U=1)



**lwzx** rD,rA,rB (X-mode, U=0)  
**lwzux** rD,rA,rB (X-mode, U=1)



if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))<sub>32:63</sub>



if X-mode then  $EA \leftarrow {}^{32}0 \parallel (a + rB)_{32:63}$   
 $rD \leftarrow {}^{32}0 \parallel MEM(EA,4)$   
 if U=1 then  $rA \leftarrow EA$

The EA is calculated as follows:

- For **lwz** and **lwzu**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **lwzx** and **lwzux**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The word addressed by the EA is loaded into rD[32–63]; rD[0–31] are cleared.

If U=1 (with update), EA is placed into rA.


If U=1 (with update), and rA=0 or rA=rD, the instruction form is invalid.

Other registers altered: None


**\_lwz** VLE User **\_lwz**

**Load Word and Zero [with Update] [Indexed]**


**e\_lwz** **rD,D(rA)** (D-mode)



**se\_lwz** **rZ,SD4(rX)** (SD4-mode)



**e\_lwzu** **rD,D8(rA)** (D8-mode)



if (RA=0 & !se\_lwz) then  $a \leftarrow {}^{32}0$  else  $a \leftarrow GPR(RA \text{ or } RX)$

if D-mode then  $EA \leftarrow (a + EXTS(D))_{32:63}$

if D8-mode then  $EA \leftarrow (a + EXTS(D8))_{32:63}$

if SD4-mode then  $EA \leftarrow (a + ({}^{26}0 \parallel SD4 \parallel {}^20))_{32:63}$

$GPR(RD \text{ or } RZ) \leftarrow MEM(EA,4)$

if e\_lwzu then  $GPR(RA) \leftarrow EA$

Let the EA be calculated as follows:

- For **e\_lwz** and **e\_lwzu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_lwz** let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 2 bits.

The word in memory addressed by the EA is loaded into bits 32–63 of GPR(rD).

If **e\_lwzu**, EA is placed into GPR(rA).

If **e\_lwzu** and rA = 0 or rA = rD, the instruction form is invalid.

Special Registers Altered: None

**mbar**

Book E	User
--------	------

**mbar**

**Memory barrier**

**mbar**

MO

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	MO					///					1	1	0	1	0	1	0	1	1	0	/					

When MO=0, **mbar** provides a memory ordering function for all memory access instructions executed by the processor executing the **mbar** instruction. Executing an **mbar** instruction ensures that all data memory accesses caused by instructions preceding the **mbar** have completed before any data memory accesses caused by any instructions after the **mbar**. This order is seen by all mechanisms.

When **mbar** (MO = 1), as defined by the EIS, **mbar** functions like **eieio** as it is defined by the Classic PowerPC architecture. It provides ordering for the effects of load and store instructions. These instructions consist of two sets, which are ordered separately. Memory accesses caused by a **dcbz** or a **dcbz** are ordered like a store. The two sets follow:

- Caching-inhibited, guarded loads and stores to memory and write-through-required stores to memory. **mbar** (MO=1) controls the order in which accesses are performed in main memory. It ensures that all applicable memory accesses caused by instructions preceding the **mbar** have completed with respect to main memory before any applicable memory accesses caused by instructions following **mbar** access main memory. It acts like a barrier that flows through the memory queues and to main memory, preventing the reordering of memory accesses across the barrier. No ordering is performed for **dcbz** if the instruction causes the system alignment error handler to be invoked.  
All accesses in this set are ordered as one set; there is not one order for guarded, caching-inhibited loads and stores and another for write-through-required stores.
- Stores to memory that are caching-allowed, write-through not required, and memory-coherency required. **mbar** (MO=1) controls the order in which accesses are performed with respect to coherent memory. It ensures that, with respect to coherent memory, applicable stores caused by instructions before the **mbar** complete before any applicable stores caused by instructions after it.

Except for **dcbz** and **dcbz**, **mbar** (MO=1) does not affect the order of cache operations (whether caused explicitly by a cache management instruction or implicitly by the cache coherency mechanism). Also, **mbar** does not affect the order of accesses in one set with respect to accesses in the other.

**mbar** (MO=1) may complete before memory accesses caused by instructions preceding it have been performed with respect to main memory or coherent memory as appropriate. **mbar** (MO=1) is intended for use in managing shared data structures, in accessing memory-mapped I/O, and in preventing load/store combining operations in main memory. For the first use, the shared data structure and the lock that protects it must be altered only by stores that are in the same set (for both cases described above). For the second use, **mbar** (MO=1) can be thought of as placing a barrier into the stream of memory accesses issued by a core, such that any given memory access appears to be on the same side of the barrier to both the core and the I/O device.

Because the core performs store operations in order to memory that is designated as both caching-inhibited and guarded, **mbar** (MO=1) is needed for such memory only when loads must be ordered with respect to stores or with respect to other loads.

Note that **mbar** (MO=1) does not connect hardware considerations to it such as multiprocessor implementations that send an **mbar** (MO=1) address-only broadcast (useful in some designs). For example, if a design has an external buffer that re-orders loads and stores for better bus efficiency, **mbar** (MO=1) broadcasts signals to that buffer that previous loads/stores (marked caching-inhibited, guarded, or write-through required) must complete before any following loads/stores (marked caching-inhibited, guarded, or write-through required).


If MO is not 0 or 1, an implementation may support the **mbar** instruction ordering a particular subset of memory accesses. An implementation may also support multiple, non-zero values of MO that each specify a different subset of memory accesses that are ordered by the **mbar** instruction. Which subsets of memory accesses are ordered and which values of MO specify these subsets is implementation-dependent. See the user’s manual for the implementation.

On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

Other registers altered: None

Programming note: **mbar** is provided to implement a pipelined memory barrier. The following sequence shows one use of **mbar** in supporting shared data, ensuring the action is completed before releasing the lock.

<b>P1</b>	<b>P2</b>
lock	. . .
read & write	. . .
mbar	. . .
free lock	. . .
. . .	lock
. . .	read & write
. . .	mbar
. . .	free lock



$$CR_{4 \times BF + 32 : 4 \times BF + 35} \leftarrow CR_{4 \times crS + 32 : 4 \times crS + 35}$$

The contents of field **crS** (bits  $4 \times crS + 32 - 4 \times crS + 35$ ) of CR are copied to field **crD** (bits  $4 \times crD + 32 - 4 \times crD + 35$ ) of CR.

Other registers altered: CR

**\_mcrf** VLE    User **\_mcrf**

**Move CR Field**

**e\_mcrf** **crD,crS**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	1	CRD	//	CRS	///						0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	/

$$CR_{4 \times CRD + 32 : 4 \times CRD + 35} \leftarrow CR_{4 \times CRS + 32 : 4 \times CRS + 35}$$

The contents of field **crS** (bits  $4 \times CRS + 32$  through  $4 \times CRS + 35$ ) of the CR are copied to field **crD** (bits  $4 \times CRD + 32$  through  $4 \times CRD + 35$ ) of the CR.

Special Registers Altered: CR

**mcrfs** Book E    User **mcrfs**

**Move to condition register from FPSCR**

**mcrfs** **crD,crS**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
1	1	1	1	1	1	1	crD	//	crS	///						0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	/

$$CR_{BF \times 4 : crD \times 4 + 3} \leftarrow FPSCR_{crS \times 4 : crS \times 4 + 3}$$

$$FPSCR_{crS \times 4 : crS \times 4 + 3} \leftarrow 0b0000$$

The contents of  $FPSCR[crS]$  are copied to CR field **crD**. All exception bits copied are cleared in the FPSCR. If the FX bit is copied, it is cleared in the FPSCR.

If  $MSR[FP]=0$ , an attempt to execute **mcrfs** causes a floating-point unavailable interrupt.

Other registers altered:

- CR field **crD**
- FX OX (if **crS**=0)
- UX ZX XX VXSNaN (if **crS**=1)
- VXISI VXIDI VXZDZ VXIMZ (if **crS**=2)
- VXVC (if **crS**=3)
- VXSOFT VXSQRT VXCVI (if **crS**=5)

**mcrxr** Book E    User **mcrxr**

**Move to condition register from integer exception register**

**mcrxr** **crD**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1	1	crD	///						1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	/

$$CR_{4 \times crD + 32 : 4 \times crD + 35} \leftarrow XER_{32 : 35}$$

$$XER_{32 : 35} \leftarrow 0b0000$$

The contents of  $XER[32-35]$  are copied to CR field **crD**.  $XER[32-35]$  are cleared.

Other registers altered: CR  $XER[32-35]$




**mfapidi**

Book E	User
--------	------

**mfapidi**

**Move from APID Indirect**

**mfapidi**                      **rD,rA**



rD " implementation-dependent value based on rA

The contents of rA are provided to any auxiliary processing extensions that may be present. A value, that is implementation-dependent and extension-dependent, is placed in rD.

Other registers altered: None

Programming note: This instruction is provided as a mechanism for software to query the presence and configuration of one or more auxiliary processing extensions. See user's manual for the implementation for details on the behavior of this instruction.


**\_mfar**

VLE	User
-----	------

**\_mfar**

**Move from Alternate Register**

**se\_mfar**                      **rX,arY**



GPR(RX) ← GPR(ARY)

For **se\_mfar**, the contents of GPR(arY) are placed into GPR(rX). arY specifies a GPR in the range R8–R23. The encoding 0000 specifies R8, 0001 specifies R9, ..., 1111 specifies R23.

Special Registers Altered: None


**mfcr**

Book E	User
--------	------

**mfcr**

**Move from condition register**

**mfcr**                      **rD**



rD ← <sup>32</sup>0 || CR

The contents of the CR are placed into rD[32–63]. Bits rD[0–31] are cleared.

Other registers altered: None


**\_mfctr**

VLE	User
-----	------

**\_mfctr**

**Move From Count Register**

**se\_mfctr**                      **rX**



GPR(RX) ← CTR

The CTR contents are placed into bits 32–63 of GPR(rX).

Special Registers Altered: None

**mfdcr** Book E | User **mfdcr**  
**Move from device control register**  
**mfdcr** rD,DCRN

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					DCRN <sub>5-9</sub>					DCRN <sub>0-4</sub>					0	1	0	1	0	0	0	0	1	1	/

rD ← DCREG(DCRN)

DCRN identifies the DCR (see the user’s manual for a list of DCRs supported by the implementation).

The contents of the designated DCR are placed into rD. For 32-bit DCRs, the contents of the DCR are placed into rD[32–63]. Bits rD[0–31] are cleared.

Execution of this instruction is restricted to supervisor mode.

Other registers altered: None

**mffs** Book E | User **mffs**  
**Move from FPSCR**  
**mffs** frD (Rc=0)  
**mffs.** frD (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	1	1	1	1	frD					///					1	0	0	1	0	0	0	1	1	1	Rc					

frD ← FPSCR

The contents of the FPSCR are placed into frD[32–63]; frD[0–31] are undefined.

If MSR[FP]=0, an attempt to execute **mffs**[.] causes a floating-point unavailable interrupt.

Other registers altered:

- CR1 ← FX || FEX || VX || OX (if Rc=1)

**\_mflr** VLE | User **\_mflr**  
**Move From Link Register**

**se\_mflr** rX

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0 0 1 0 0 0					RX				

GPR(RX) ← LR

The LR contents are placed into bits 32–63 of GPR(rX).

Special Registers Altered: None

**mfmsr**

Book E	User
--------	------

**mfmsr**

**Move from machine state register**

**mfmsr**

**rD**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					///					0	0	0	1	0	1	0	0	1	1	/					

$$rD \leftarrow {}^{32}0 \parallel \text{MSR}$$

The contents of the MSR are placed into rD[32–63]. Bits rD[0–31] are cleared.

Execution of this instruction is restricted to supervisor mode.

Other registers altered: None

**mfpmr**

Book E	User
--------	------

**mfpmr**

**Move from Performance Monitor Register**

**mfpmr**

**rD,PMRN**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					PMRN <sub>5–9</sub>				PMRN <sub>0–4</sub>				0	1	0	1	0	0	1	1	1	0	0		

$$\text{GPR}(rD) \leftarrow \text{PMREG}(\text{PMRN})$$

PMRN denotes a performance monitor register. [Section 3.16: Performance monitor registers \(PMRs\)](#), lists supported performance monitor registers.

The contents of the designated performance monitor register are placed into GPR[rD].

When MSR[PR] = 1, specifying a performance monitor register that is not implemented and is not privileged (PMRN[5] = 0) results in an illegal instruction exception-type program interrupt. When MSR[PR] = 1, specifying a performance monitor register that is privileged (PMRN[5] = 1) results in a privileged instruction exception-type program interrupt. When MSR[PR] = 0, specifying an unimplemented performance monitor register is boundedly undefined.

Other registers altered: None

**mfspr**

Book E	User
--------	------

**mfspr**

**Move from special purpose register**

**mfspr**

**rD,SPRN**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					SPRN[5–9]				SPRN[0–4]				0	1	0	1	0	1	0	0	1	1	/		

$$rD \leftarrow \text{SPREG}(\text{SPRN})$$

SPRN denotes an SPR (see [Section 3.18: Book E SPR model](#)).



The contents of the designated SPR are placed into rD. For 32-bit SPRs, the contents of the SPR are placed into rD[32–63]. Bits rD[0–31] are cleared.

**Table 207. Effect of SPRN[5] and MSR[PR]**

SPRN[5]	MSR[PR]	SPRN class	Result
0	1	Defined	If not implemented, illegal instruction exception If implemented, as defined in Book E
0	1	Allocated	If not implemented, illegal instruction exception If implemented, as defined in user's manual
0	1	Preserved	If not implemented, illegal instruction exception If implemented, as defined in PowerPC Architecture
0	1	Reserved	Illegal instruction exception
1	1	—	Privileged exception
—	0	Defined	If not implemented, boundedly undefined If implemented, as defined in Book E
—	0	Allocated	If not implemented, boundedly undefined If implemented, as defined in user's manual
—	0	Preserved	If not implemented, boundedly undefined If implemented, as defined in PowerPC Architecture
—	0	Reserved	Boundedly undefined

Execution of this instruction specifying a defined and privileged SPR (SPRN[5]=1) when MSR[PR]=1 results in a privileged instruction exception-type program interrupt.


Other registers altered: None



**Move Register**

**se\_mr**

rX,rY



$$GPR(RX) \leftarrow GPR(RY)$$


For **se\_mr**, the contents of GPR(rY) are placed into GPR(rX).

Special Registers Altered: None



**Memory synchronize**

**msync**



The **msync** instruction provides an ordering function for the effects of all instructions executed by the processor executing the **msync**. Executing **msync** ensures that all instructions preceding the **msync** have completed before **msync** completes and that no subsequent instructions are initiated until after the **msync** completes. It also creates a memory barrier (see [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#)), which orders the memory accesses associated with these instructions.

The **msync** may not complete before memory accesses associated with instructions preceding **msync** have been performed.


On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

**msync** is execution synchronizing. (See [Section 4.2.3.7: Execution synchronization.](#))

Other registers altered: None

Programming notes:


- **msync** can be used to ensure that all stores into a data structure, caused by store instructions executed in a critical section of a program, are performed with respect to another processor before the store that releases the lock is performed with respect to that processor.  
The functions performed by the **msync** may take a significant amount of time to complete, so indiscriminate use of this instruction may adversely affect performance. The Memory Barrier (**mbar**) instruction may be more appropriate than **msync** for many cases.
- **msync** replaces the **sync** instruction; it uses the same opcode as **sync** such that PowerPC applications calling for **sync** invoke the **msync** when executed on an Book E implementation. The functionality of **msync** is identical to **sync** except that **msync** also does not complete until all previous memory accesses complete. **mbar** is provided in the Book E for those occasions when only ordering of memory accesses is required without execution synchronization.



$$GPR(ARX) \leftarrow GPR(RY)$$

For **se\_mtar**, the contents of GPR(rY) are placed into GPR(arX). arX specifies a GPR in the range R8–R23. The encoding 0000 specifies R8, 0001 specifies R9,..., 1111 specifies R23.

Special Registers Altered: None



```

i ← 0
do while i < 8
    if CRMi=1 then CR4×i+32:4×i+35 ← rS4×i+32:4×i+35
    i ← i+1


```

The contents of rS[32–63] are placed into the CR under control of the field mask specified by CRM. The field mask identifies the 4-bit fields affected. Let i be an integer in the range 0–7. If CRM<sub>i</sub> = 1, CR field i (CR bits 4×i+32 through 4×i+35) is set to the contents of the corresponding field of rS[32–63].

Other registers altered: CR fields selected by mask




### Move To Count Register



```
CTR ← GPR(RX)
```


The contents of bits 32–63 of GPR(rX) are placed into the CTR.

Special Registers Altered: CTR



### Move to device control register

**mtdcr**    DCRN,rS




```
DCREG(DCRN) ← rS
```

DCRN identifies the DCR (see user’s manual for a list of DCRs supported by the implementation).

The contents of rS are placed into the designated DCR. For 32-bit DCRs, rS[32–63] are placed into the DCR.


Execution of this instruction is restricted to supervisor mode.

Other registers altered: See the user’s manual for the implementation



### Move to FPSCR Bit 0

**mtfsb0**    **crbD**    (Rc=0)  
**mtfsb0.**    **crbD**    (Rc=1)



FPSCR[BT] ← 0b0

FPSCR[BT] is cleared.

If MSR[FP]=0, an attempt to execute **mtfsb0**[.] causes a floating-point unavailable interrupt.

Other registers altered:

- FPSCR[BT]  
CR1 ← FX || FEX || VX || OX (if Rc=1)

Programming note: Bits 1 and 2 (FEX and VX) cannot be explicitly reset.

**mtfsb1** Book E User **mtfsb1**

Move to FPSCR Bit 1

**mtfsb1** **crbD** (Rc=0)  
**mtfsb1.** **crbD** (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	1	1	1	1	<b>crbD</b>					///					0	0	0	0	1	0	0	1	1	0	Rc					

FPSCR[BT] ← 0b1

FPSCR[BT] is set.

If MSR[FP]=0, an attempt to execute **mtfsb1**[.] causes a floating-point unavailable interrupt.

Other registers altered:

- FPSCR[BT,FX]  
CR1 ← FX || FEX || VX || OX (if Rc=1)

Programming note: Bits 1 and 2 (FEX and VX) cannot be explicitly set.

**mtfsf** Book E User **mtfsf**

Move to FPSCR fields

**mtfsf** **FM,frB** (Rc=0)  
**mtfsf.** **FM,frB** (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	1	1	1	1	/	<b>FM</b>					/	<b>frB</b>					1	0	1	1	0	0	0	1	1	1	Rc			

```

i ← 0
do while i < 8
    if FMi = 1 then FPSCR4×i:4×i+3 ← frB4×i:4×i+3
    i ← i + 1
    
```

The contents of **frB**[32–63] are placed into the FPSCR under control of the field mask specified by **FM**. The field mask identifies the 4-bit fields affected. Let *i* be an integer in the range 0–7. If **FM**<sub>*i*</sub> = 1, FPSCR field *i* (FPSCR bits 4×*i* through 4×*i*+3) is set to the contents of the corresponding field of the low-order 32 bits of **frB**.

FPSCR[FX] is altered only if **FM**<sub>0</sub> = 1.

If MSR[FP]=0, an attempt to execute **mtfsf**[.] causes a floating-point unavailable interrupt.


Other registers altered:

- FPSCR fields selected by mask  
 $CR1 \leftarrow FX \parallel FEX \parallel VX \parallel OX$  (if  $Rc=1$ )

Programming notes:

- Updating fewer than all eight fields of the FPSCR may have substantially poorer performance on some implementations than updating all the fields.
- When  $FPSCR[0-3]$  is specified, bits 0 (FX) and 3 (OX) are set to the values of  $(frB)_{32}$  and  $(frB)_{35}$  (that is, even if this instruction causes OX to change from 0 to 1, FX is set from  $(frB)_{32}$  and not by the usual rule that FX is set when an exception bit changes from 0 to 1). Bits 1 and 2 (FEX and VX) are set according to the usual rule (see [Table 10: FPSCR field descriptions](#)) and not from  $(frB)_{33-34}$ .

<b>mtfsfi</b>	Book E	User	<b>mtfsfi</b>
<b>Move to FPSCR field immediate</b>			
<b>mtfsfi</b>	<b>crD,UIMM</b>		( $Rc=0$ )
<b>mtfsfi.</b>	<b>crD,UIMM</b>		( $Rc=1$ )



$$FPSCR_{BF \times 4: crD \times 4+3} \leftarrow UIMM$$

The value of the UIMM field is placed into  $FPSCR[crD]$ .

$FPSCR[FX]$  is altered only if  $crD = 0$ .

If  $MSR[FP]=0$ , an attempt to execute **mtfsfi.** causes a floating-point unavailable interrupt.

Other registers altered:

- $FPSCR[crD]$   
 $CR1 \leftarrow FX \parallel FEX \parallel VX \parallel OX$  (if  $Rc=1$ )

Programming note: When  $FPSCR[0-3]$  is specified, bits 0 (FX) and 3 (OX) are set to the values of U0 and U3 (that is, even if this instruction causes OX to change from 0 to 1, FX is set from U0 and not by the usual rule that FX is set when an exception bit changes from 0 to 1). Bits 1 and 2 (FEX and VX) are set according to the usual rule (see [Table 10: FPSCR field descriptions](#)), and not from U1–2.

<b>_mtlr</b>	VLE	User	<b>_mtlr</b>
<b>Move To Link Register</b>			
<b>se_mtlr</b>	<b>rX</b>		
	0	5 6	11 12 15
	0 0 0 0 0 0	0 0 1 0 0 1	RX

$$LR \leftarrow GPR(RX)$$

The contents of bits 32–63 of  $GPR(rX)$  are placed into the LR.

Special Registers Altered: LR

**mtmsr**


Book E	Supervisor
--------	------------

**mtmsr**

**Move to machine state register**

**mtmsr**

**rS**



$$MSR \leftarrow rS_{32:63}$$

The contents of rS[32–63] are placed into the MSR.

Execution of this instruction is restricted to supervisor mode.

Execution of this instruction is execution synchronizing. See [Section 4.2.3.7: Execution synchronization](#).

In addition, changes to the EE or CE bits are effective as soon as the instruction completes. Thus if MSR[EE]=0 and an external interrupt is pending, executing an **mtmsr** that sets MSR[EE] causes the external interrupt to be taken before the next instruction is executed, if no higher priority exception exists. Likewise, if MSR[CE]=0 and a critical input interrupt is pending, executing an **mtmsr** that sets MSR[CE] causes the critical input interrupt to be taken before the next instruction is executed if no higher priority exception exists.

Other registers altered: MSR

Programming note: For a discussion of software synchronization requirements when altering certain MSR bits, refer to [Section 3.18.2: Synchronization requirements for SPRs](#).

**mtpmr**


Performance Monitor APU	User/Supervisor
-------------------------	-----------------

**mtpmr**

**Move To Performance Monitor Register**

**mtpmr**

**PMRN,rS**



$$PMREG(PMRN) \leftarrow GPR(RS)$$

PMRN denotes a performance monitor register. [Section 3.16: Performance monitor registers \(PMRs\)](#), lists supported performance monitor registers).

The contents of GPR[rS] are placed into the designated performance monitor register.

When MSR[PR] = 1, specifying a performance monitor register that is not implemented and is not privileged (PMRN[5] = 0) results in an illegal instruction exception-type program interrupt. When MSR[PR] = 1, specifying a performance monitor register that is privileged (PMRN[5] = 1) results in a privileged instruction exception-type program interrupt. When MSR[PR] = 0, specifying a unimplemented performance monitor register is boundedly undefined.

Other registers altered: None



**mtspr**

Book E	User/Supervisor
--------	-----------------

**mtspr**

**Move to special purpose register**

**mtspr** SPRN,rS

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rS					SPRN[5–9]					SPRN[0–4]					0	1	1	1	0	1	0	0	1	1	/

$$\text{SPREG}(\text{SPRN}) \leftarrow \text{rS}$$

SPRN denotes an SPR (see [Section 3.18: Book E SPR model](#), and the user’s manual of the implementation for a list of all SPRs that are implemented).

The contents of rS are placed into the designated SPR. For 32-bit SPRs, the contents of rS[32–63] are placed into the SPR.

When MSR[PR]=1, specifying an SPR that is not implemented and is not privileged (SPRN[5]=0) results in an illegal instruction exception-type program interrupt. When MSR[PR]=1, specifying an SPR that is privileged (SPRN[5]=1) results in a privileged instruction exception-type program interrupt. When MSR[PR]=0, specifying an SPR that is not implemented is boundedly undefined.

Other registers altered: See [Section 3.18: Book E SPR model](#), or the user’s manual for the implementation.

Programming note: For a discussion of software synchronization requirements when altering certain SPRs, please refer to [Section 3.18.2: Synchronization requirements for SPRs](#).

**mulhw**

Book E	User
--------	------

**mulhw**

**Multiply high word**

**mulhw** rD,rA,rB (Rc=0)  
**mulhw.** rD,rA,rB (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	rD					rA					rB					/	0	0	1	0	0	1	0	1	1	Rc

$$\text{prod}_{0:63} \leftarrow \text{rA}_{32:63} \times \text{rB}_{32:63}$$

if Rc=1 then do

$$\text{LT} \leftarrow \text{prod}_{0:31} < 0$$

$$\text{GT} \leftarrow \text{prod}_{0:31} > 0$$

$$\text{EQ} \leftarrow \text{prod}_{0:31} = 0$$

$$\text{CR0} \leftarrow \text{LT} \parallel \text{GT} \parallel \text{EQ} \parallel \text{SO}$$

$$\text{rD}_{32:63} \leftarrow \text{prod}_{0:31}$$

$$\text{rD}_{0:31} \leftarrow \text{undefined}$$

Bits 0–31 of the 64-bit product of the contents of rA[32–63] and the contents of rB[32–63] are placed into rD[32–63]. Bits rD[0–31] are undefined.

Both operands and the product are interpreted as signed integers.

Other registers altered: CR0 (if Rc=1)




**mulhwu**

Book E	User
--------	------

**mulhwu**

**Multiply high word unsigned**

**mulhwu**  $rD, rA, rB$  (Rc=0)  
**mulhwu.**  $rD, rA, rB$  (Rc=1)



```

prod0:63 ← rA32:63 × rB32:63
if Rc=1 then do
    LT ← prod0:31 < 0
    GT ← prod0:31 > 0
    EQ ← prod0:31 = 0
    CR0 ← LT || GT || EQ || SO

    rD32:63 ← prod0:31
    rD0:31 ← undefined
    
```

Bits 0–31 of the 64-bit product the contents of rA[32–63] and the contents of rB[32–63] are placed into rD[32–63]. Bits rD[0–31] are undefined.

Both operands and the product are interpreted as unsigned integers, except that if Rc=1 the first three bits of CR field 0 are set by signed comparison of the result to zero.

Other registers altered: CR0 (if Rc=1)


**mulli**

Book E	User
--------	------

**mulli**

**Multiply low immediate**

**mulli**  $rD, rA, SIMM$



```

prod0:127 ← rA × EXTS(SIMM)
rD ← prod64:127
    
```

Bits 64–127 of the 128-bit product of the contents of rA and the sign-extended value of the SIMM field are placed into rD.

Both operands and the product are interpreted as signed integers.

Other registers altered: None

Programming notes:

- For **mulli**, the low-order 64 bits of the product are independent of whether the operands are regarded as signed or unsigned 64-bit integers.
- For **mulli** and **mullw**, bits 32–63 of the product are independent of whether the operands are regarded as signed or unsigned 32-bit integers.



**\_mullix**

VLE
-----

User
------


**\_mullix**

**Multiply Low [2 operand] Immediate**

**e\_mulli**

**rD,rA,SCI8**

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31



imm ← SCI8(F,SCL,UI8)  
 prod<sub>0:63</sub> ← GPR(RA) × imm  
 GPR(RD) ← prod<sub>32:63</sub>

Bits 32–63 of the 64-bit product of the contents of GPR(rA) and the value of SCI8 are placed into GPR(rD).


Both operands and the product are interpreted as signed integers.

Special Registers Altered: None

**e\_mull2i**

**rA,SI**

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31



prod<sub>0:63</sub> ← GPR(RA) × EXTS(SI<sub>0:4</sub> || SI<sub>5:15</sub>)  
 GPR(RA) ← prod<sub>32:63</sub>

Bits 32–63 of the 64-bit product of the contents of GPR(rA) and the sign-extended value of the SI field are placed into GPR(rA).

Both operands and the product are interpreted as signed integers.

Special Registers Altered: None

**mullw**

Book E
--------


User
------

**mullw**

**Multiply low word**

<b>mullw</b>	<b>rD,rA,rB</b>	(OE=0, Rc=0)
<b>mullw.</b>	<b>rD,rA,rB</b>	(OE=0, Rc=1)
<b>mullwo</b>	<b>rD,rA,rB</b>	(OE=1, Rc=0)
<b>mullwo.</b>	<b>rD,rA,rB</b>	(OE=1, Rc=1)

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31



prod<sub>0:63</sub> ← rA<sub>32:63</sub> × rB<sub>32:63</sub>  
 if OE=1 then do  
     OV ← (prod<sub>0:31</sub> ≠ <sup>32</sup>0) & (prod<sub>0:31</sub> ≠ <sup>32</sup>1)  
     SO ← SO | OV  
 if Rc=1 then do  
     LT ← prod<sub>32:63</sub> < 0  
     GT ← prod<sub>32:63</sub> > 0  
     EQ ← prod<sub>32:63</sub> = 0  
     CR0 ← LT || GT || EQ || SO  
 rD ← prod<sub>0:63</sub>



The 64-bit product of the contents of rA[32–63] and the contents of rB[32–63] is placed into rD.

If OE=1, OV is set if the product cannot be represented in 32 bits.

Both operands and the product are interpreted as signed integers.

Other registers altered:

- CR0 (if Rc=1)  
SO OV (if OE=1)

Programming notes:


- Bits 32–63 of the product are independent of whether the operands are regarded as signed or unsigned 32-bit integers.

**\_mullwx**

VLE	User
-----	------

**\_mullwx**

**Multiply Low Word**



$$prod_{0:63} \leftarrow GPR(RX)_{32:63} \times GPR(RY)_{32:63}$$

$$GPR(RX) \leftarrow prod_{32:63}$$

Bits 32–63 of the 64-bit product of the contents of bits 32–63 of GPR(rX) and the contents of bits 32–63 of GPR(rY) is placed into GPR(rX).

Special Registers Altered: None


**nand**

Book E	User
--------	------

**nand**

**NAND**

**nand** **rA,rS,rB** (Rc=0)  
**nand.** **rA,rS,rB** (Rc=1)



$$result_{0:63} \leftarrow \neg(rS \& rB)$$

if Rc=1 then do

$$LT \leftarrow result_{32:63} < 0$$

$$GT \leftarrow result_{32:63} > 0$$

$$EQ \leftarrow result_{32:63} = 0$$

$$CR0 \leftarrow LT \parallel GT \parallel EQ \parallel SO$$


$$rA \leftarrow result$$

The contents of rS are ANDed with the contents of rB and the one's complement of the result is placed into rA.

Other registers altered: CR0 (if Rc=1)



<b>neg</b>	Book E	User	<b>neg</b>
<b>Negate</b>			
<b>neg</b>	rD,rA		(OE=0, Rc=0)
<b>neg.</b>	rD,rA		(OE=0, Rc=1)
<b>nego</b>	rD,rA		(OE=1, Rc=0)
<b>nego.</b>	rD,rA		(OE=1, Rc=1)



```

carry0:63 ← Carry(−rA + 1)
sum0:63 ← −rA + 1
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← sum
    
```


The sum of the one’s complement of the contents of rA and 1 is placed into rD.

If rA contains the most negative 64-bit number (0x8000\_0000\_0000\_0000), the result is the most negative number. Similarly, if rA[32–63] contain the most negative 32-bit number (0x8000\_0000), bits 32–63 of the result contain the most negative 32-bit number and, if OE=1, OV is set.

Other registers altered:

- CR0 (if Rc=1)
- SO OV (if OE=1)

<b>_negx</b>	VLE	User	<b>_negx</b>
<b>Negate</b>			
<b>se_neg</b>	rX		



```


result32:63 ← −GPR(RX)+ 1
GPR(RX) ← result32:63
    
```

The sum of the one’s complement of the contents of GPR(rX) and 1 is placed into GPR(rX).

If bits 32–63 of GPR(rX) contain the most negative 32-bit number (0x8000\_0000), bits 32–63 of the result contain the most negative 32-bit number

Special Registers Altered: None

**nor** Book E User **nor**  
**NOR**  
**nor** **rA,rS,rB** (Rc=0)  
**nor.** **rA,rS,rB** (Rc=1)




result<sub>0:63</sub> ← ¬(rS | rB)  
 if Rc=1 then do  
     LT ← result<sub>32:63</sub> < 0  
     GT ← result<sub>32:63</sub> > 0  
     EQ ← result<sub>32:63</sub> = 0  
     CR0 ← LT || GT || EQ || SO  
 rA ← result

The contents of rS are ORed with the contents of rB and the one's complement of the result is placed into rA.

Other registers altered: CR0 (if Rc=1)

**\_notx** VLE User **\_notx**  
**NOT**  
**se\_notrX**




result<sub>32:63</sub> ← ¬GPR(RX)  
 GPR(RX) ← result<sub>32:63</sub>


The contents of GPR(rX) are inverted.

Special Registers Altered: None


**or** Book E User **or**  
**OR [Immediate [shifted] | with complement]**  
**or** **rA,rS,rB** (Rc=0)  
**or.** **rA,rS,rB** (Rc=1)



**ori** **rA,rS,UIMM** (S=0, Rc=0)  
**oris** **rA,rS,UIMM** (S=1, Rc=0)



**orc** rA,rS,rB (Rc=0)  
**orc.** rA,rS,rB (Rc=1)







if 'ori' then  $b \leftarrow 48_0 \parallel \text{UIMM}$   
 if 'oris' then  $b \leftarrow 32_0 \parallel \text{UIMM} \parallel 16_0$   
 if 'or[.]' then  $b \leftarrow rB$   
 if 'orc[.]' then  $b \leftarrow \neg rB$   
 $\text{result}_{0:63} \leftarrow rS \mid b$   
 if Rc=1 then do  
      $LT \leftarrow \text{result}_{32:63} < 0$   
      $GT \leftarrow \text{result}_{32:63} > 0$   
      $EQ \leftarrow \text{result}_{32:63} = 0$   
      $CR0 \leftarrow LT \parallel GT \parallel EQ \parallel SO$   
 $rA \leftarrow \text{result}$

For **ori**, the contents of rS are ORed with  $48_0 \parallel \text{UIMM}$ .  
 For **oris**, the contents of rS are ORed with  $32_0 \parallel \text{UIMM} \parallel 16_0$ .  
 For **or[.]**, the contents of rS are ORed with the contents of rB.  
 For **orc[.]**, the contents of rS are ORed with the one's complement of the contents of rB.  
 The result is placed into rA.  
 The preferred no-op is **ori 0,0,0**  
 Other registers altered: CR0 (if Rc=1)

**\_orx**

VLE	User
-----	------

**\_orx**  
**OR [2 operand] [Immediate | with Complement] [Shifted][and Record]**



if 'e\_ori[.]' then  $b \leftarrow \text{SCI8}(F,SCL,UI8)$



```

if 'e_or2i' then b ← 160 || UI0:4 || UI5:15
if 'e_or2is' then b ← UI0:4 || UI5:15 || 160
if 'se_or' then b ← GPR(RB)

result0:63 ← GPR(RS or RD or RX) | b

if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
    GPR(RA or RD or RX) ← result
    
```

For **e\_ori[.]**, the contents of GPR(rS) are ORed with the value of SCI8.

For **e\_or2i**, the contents of GPR(rD) are ORed with <sup>16</sup>0 || UI.

For **e\_or2is**, the contents of GPR(rD) are ORed with UI || <sup>16</sup>0.

For **se\_or**, the contents of GPR(rX) are ORed with the contents of GPR(rY).

The result is placed into GPR(rA or rX).

The preferred 'no-op' (an instruction that does nothing) is:

```
e_ori 0,0,0
```

Special Registers Altered: CR0 (if Rc = 1)


**rfci**

Book E	Supervisor
--------	------------

**rfci**

**Return from critical interrupt**

**rfci**



```

MSR ← CSRR1
NIA ← CSRR0[0:61] || 0b00
    
```

The **rfci** instruction is used to return from a critical class interrupt, or as a means of establishing a new context and synchronizing on that new context simultaneously.

The contents of CSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address CSRR0[0–61]||0b00. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case, the value placed into SRR0 or CSRR0 by the interrupt processing mechanism is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in CSRR0 at the time of the execution of the **rfci**).

Execution of this instruction is restricted to supervisor mode.

Execution of this instruction is context synchronizing. See [Section 4.2.3.6: Context synchronization](#).




Other registers altered: MSR

Programming note: In addition to Branch to LR (**bclr[l]**) and Branch to CTR (**bcctr[l]**) instructions, **rfi** and **rfci** allow software to branch to any valid 64-bit address by using the respective 64-bit SRR0 and CSRR0.

**\_rfci** VLE Supervisor **\_rfci**

**Return From Critical Interrupt**

**se\_rfc**



MSR ← CSRR1  
 NIA ← CSRR0<sub>0:62</sub> || 0b0

The **se\_rfc** instruction is used to return from a critical class interrupt, or as a means of establishing a new context and synchronizing on that new context simultaneously.

The contents of CSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address CSRR0[32–62]||0b0. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism (see Book E) is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in CSRR0 at the time of the execution of the **se\_rfc**).

Execution of this instruction is privileged and restricted to supervisor mode.


Execution of this instruction is context synchronizing.

Special Registers Altered: MSR

**rfdi** Debug APU Supervisor **rfdi**

**Return from debug interrupt**

**rfdi**



if Mode32 then m ← 32  
 if Mode64 then m ← 0  
 MSR ← DSRR1  
 NIA ← <sup>m</sup>0 || DSRR0<sub>m:61</sub> || 0b00

The **rfdi** instruction is used to return from a debug interrupt, or as a means of establishing a new context and synchronizing on that new context simultaneously.

The contents of DSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address DSRR0[0–61]||0b00. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0, CSRR0, or DSRR0 by the interrupt processing mechanism is the address of the instruction that would have been executed next



had the interrupt not occurred (that is, the address in DSRR0 at the time of the execution of the **rfdi**).

Execution of this instruction is privileged and restricted to supervisor mode.

Execution of this instruction is context synchronizing.

Other registers altered:

- MSR set as described above.


**rfi**

Book E	Supervisor
--------	------------

**rfi**

**Return from interrupt**

**rfi**



MSR ← SRR1  
 NIA ← SRR0[0:61] || 0b00

The **rfi** instruction is used to return from a non-critical class interrupt, or as a means of simultaneously establishing a new context and synchronizing on that new context.

The contents of SRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address SRR0[0–61]||0b00. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in SRR0 at the time of the execution of the **rfi**).

Execution of this instruction is restricted to supervisor mode.

Execution of this instruction is context synchronizing. See [Section 4.2.3.6: Context synchronization](#).

Other registers altered: MSR


**\_rfi**

VLE	Supervisor
-----	------------

**\_rfi**

**Return From Interrupt**

**se\_rfi**



MSR ← SRR1  
 NIA ← SRR0<sub>0:62</sub> || 0b0

The **se\_rfi** instruction is used to return from a non-critical class interrupt, or as a means of simultaneously establishing a new context and synchronizing on that new context.

The contents of SRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched under control of the new MSR value





from the address  $SRR0[32-62]||0b0$ . If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism (see Book E) is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in SRR0 at the time of the execution of the **se\_rfi**).

Execution of this instruction is privileged and restricted to supervisor mode.


Execution of this instruction is context synchronizing.

Special Registers Altered: MSR

**rfmci** Machine Check APU Supervisor **rfmci**

**Return from Machine Check Interrupt**

**rfmci**



MSR ← MCSRR1  
 NIA ← MCSRR0<sub>0:61</sub> || 0b00

The **rfmci** instruction is used to return from a machine check interrupt, or as a means of simultaneously establishing a new context and synchronizing on that new context.

The contents of machine check save/restore register 1 (MCSRR1) are placed into the MSR. If the new MSR value does not enable any pending exceptions, the next instruction is fetched, under control of the new MSR value from the address  $MCSRR0[32-61]||0b00$ . If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in MCSRR0 at the time of the execution of **rfi** or **rfci**).


Execution of this instruction is privileged and context synchronizing.

Special registers altered: MSR


**\_rlw** VLE User **\_rlw**

**Rotate Left Word [Immediate]**

**e\_rlwi** (Rc = 0)  
**e\_rlwi.** (Rc = 1)



**e\_rlwi** (Rc = 0)  
**e\_rlwi.** (Rc = 1)



```

if 'e_rlw[.]' then n ← GPR(RB)59:63
else          n ← SH
result32:63 ← ROTL32(GPR(RS)32:63,n)
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
        GPR(RA) ← result32:63

```

If **e\_rlw[.]**, let the shift count *n* be the contents of bits 59–63 of GPR(**rB**).

If **e\_rlwi[.]**, let the shift count *n* be SH.

The contents of GPR(**rS**) are rotated<sub>32</sub> left *n* bits. The rotated data is placed into GPR(**rA**).

Special Registers Altered: CR0 (if Rc = 1)


**rlwimi**

Book E	User
--------	------

**rlwimi**

**Rotate left word immediate then mask insert**

**rlwimi**                    **rA,rS,SH,MB,ME**                    (Rc=0)  
**rlwimi.**                    **rA,rS,SH,MB,ME**                    (Rc=1)



```

n ← SH
b ← MB+32
e ← ME+32
r ← (rS32:63,n)
m ← MASK(b,e)
result0:63 ← r&m | rA&~m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO

rA ← result0:63

```

The shift count *n* is the value SH.

The contents of **rS** are rotated<sub>32</sub> left *n* bits. A mask is generated having 1 bits from bit MB+32 through bit ME+32 and 0 bits elsewhere. The rotated data is inserted into **rA** under control of the generated mask. (If a mask bit is 1 the associated bit of the rotated data is placed into the target register, and if the mask bit is 0 the associated bit in the target register remains unchanged.)

Other registers altered: CR0 (if Rc=1)


Programming note: Uses for **rlwimi**[.]:

- To insert a *k*-bit field that is left-justified in **rS**[32–63], into **rA**[32–63] starting at bit position *j*, by setting **SH**=64-*j*, **MB**=*j*-32, and **ME**=(*j*+*k*)-33.
- To insert an *k*-bit field that is right-justified in **rS**[32–63], into **rA**[32–63] starting at bit position *j*, by setting **SH**=64-(*j*+*k*), **MB**=*j*-32, and **ME**=(*j*+*k*)-33.

**\_rlwimi** VLE User **\_rlwimi**

**Rotate Left Word Immediate then Mask Insert**

**e\_rlwimi**    **rA,rS,SH,MB,ME**



```

n ← SH
b ← MB+32
e ← ME+32
r ← ROTL32(GPR(RS)32:63,n)
m ← MASK(b,e)
result32:63 ← r&m | GPR(RA)&¬m
GPR(RA) ← result32:63
    
```

Let the shift count *n* be the value **SH**.


The contents of **GPR(rS)** are rotated<sub>32</sub> left *n* bits. A mask is generated having 1 bits from bit **MB**+32 through bit **ME**+32 and 0 bits elsewhere. The rotated data are inserted into **GPR(rA)** under control of the generated mask (if a mask bit is 1 the associated bit of the rotated data is placed into the target register, and if the mask bit is 0 the associated bit in the target register remains unchanged).

Special Registers Altered: None

**\_rlwinm** VLE User **\_rlwinm**

**Rotate Left Word Immediate then AND with Mask**

**e\_rlwinm**    **rA,rS,SH,MB,ME**



```

n ← SH
b ← MB+32
e ← ME+32
r ← ROTL32(GPR(RS)32:63,n)
m ← MASK(b,e)
result32:63 ← r & m
GPR(RA) ← result32:63
    
```




Let the shift count  $n$  be SH.


The contents of GPR( $rS$ ) are rotated<sub>32</sub> left  $n$  bits. A mask is generated having 1 bits from bit MB+32 through bit ME+32 and 0 bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into GPR( $rA$ ).

Special Registers Altered: None

<b>rlwnm</b>	Book E	User	<b>rlwnm</b>
<b>Rotate left word [immediate] then AND with mask</b>			
<b>rlwnm</b>	$rA, rS, rB, MB, ME$		(Rc=0)
<b>rlwnm.</b>	$rA, rS, rB, MB, ME$		(Rc=1)



<b>rlwinm</b>	$rA, rS, SH, MB, ME$	(Rc=0)
<b>rlwinm.</b>	$rA, rS, SH, MB, ME$	(Rc=1)



```

if 'rlwnm[.]' then n ← rB59:63
else n ← SH
b ← MB+32
e ← ME+32
r ← (rS32-63,n)
m ← MASK(b,e)
result0:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
rA ← result0:63
    
```

If **rlwnm[.]**, the shift count,  $n$ , is the contents of  $rB[59-63]$ . If **rlwinm[.]**,  $n$  is SH. The  $rS$  contents are rotated<sub>32</sub> left  $n$  bits. The mask has 1s from bit MB+32 through bit ME+32 and 0s elsewhere. The rotated data is ANDed with the mask and the result is placed into  $rA$ .

Other registers altered: CR0 (if Rc=1)

Uses for <i>rlwnm</i> [.]	Uses for <i>rlwinm</i> [.]
To extract a <i>k</i> -bit field starting at bit position <i>j</i> in <i>rS</i> [32–63], right-justified into <i>rA</i> [32–63] (clearing the remaining 32– <i>k</i> bits of <i>rA</i> [32–63])...	...by setting <i>SH=j+k-32</i> , <i>MB=32-k</i> , and <i>ME=31</i> .
...by setting <i>rB</i> [59–63]= <i>j+k-32</i> , <i>MB=32-k</i> , and <i>ME=31</i> .	To extract a <i>k</i> -bit field that starts at bit position <i>j</i> in <i>rS</i> [32–63], left-justified into <i>rA</i> [32–63] (clearing the remaining 32– <i>k</i> bits of <i>rA</i> [32–63])...
...by setting <i>rB</i> [59–63]= <i>j-32</i> , <i>MB=0</i> , and <i>ME=k-1</i> .	...by setting <i>SH=j-32</i> , <i>MB=0</i> , and <i>ME=k-1</i> .
To rotate the contents of bits 32–63 of a register left by <i>k</i> bits...	
...setting <i>rB</i> [59–63]= <i>k</i> , <i>MB=0</i> , and <i>ME=31</i> .	...setting <i>SH=k</i> , <i>MB=0</i> , and <i>ME=31</i> .
To rotate the contents of bits 32–63 of a register right by <i>k</i> bits...	
...by setting <i>rB</i> [59–63] = <i>32-k</i> , <i>MB=0</i> , and <i>ME=31</i> .	...by setting <i>SH=32-k</i> , <i>MB=0</i> , and <i>ME=31</i> .
	To shift the contents of bits 32–63 of a register right by <i>k</i> bits, by setting <i>SH=32-k</i> , <i>MB=k</i> , and <i>ME=31</i> . To clear the high-order <i>j</i> bits of the contents of bits 32–63 of a register and then shift the result left by <i>k</i> bits, by setting <i>SH=k</i> , <i>MB=j-k</i> and <i>ME=31-k</i> . To clear the low-order <i>k</i> bits of bits 32–63 of a register, by setting <i>SH=0</i> , <i>MB=0</i> , and <i>ME=31-k</i> .
For the uses given above, bits <i>rA</i> [0–31] are cleared.	


**sc**

Book E	Supervisor
--------	------------

**sc**

**System call**

**sc**



```

SRR1 ← MSR
SRR0 ← CIA+4
NIA ← EVPR[0:47] || IVOR8[48-59] || 0b0000
MSR[WE,EE,PR,IS,DS,FP,FE0,FE1] ← 0b0000_0000
    
```

**sc** is used to request a system service. A system call interrupt is generated. The MSR contents are copied into SRR1 and the address of the instruction after the **sc** instruction is placed into SRR0.

MSR[WE,EE,PR,IS,DS,FP,FE0,FE1] are cleared.

The interrupt causes the next instruction to be fetched from the address IVPR[0–47]||IVOR8[48-59]||0b0000.

**sc** is context synchronizing. See [Section 4.2.3.6: Context synchronization](#).

Other registers altered: SRR0 SRR1 MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]




**slw** Book E   User **slw**

**Shift left word**

**slw** rA,rS,rB (Rc=0)

**slw.** rA,rS,rB (Rc=1)



```

n ← rB59:63
r ← ROTL32(rS32:63,n)
if rB58=0 then m ← MASK(32,63-n)
else m ← 640
result0:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
rA ← result0:63
    
```

The shift count n is the value specified by the contents of rB[58–63].

The contents of rS[32–63] are shifted left n bits. Bits shifted out of position 32 are lost. Zeros are supplied to the vacated positions on the right. The 32-bit result is placed into rA[32–63]. Bits rA[0–31] are cleared.

Shift amounts from 32 to 63 give a zero result.




Other registers altered: CR0 (if Rc=1)

**\_slwx** VLE   User **\_slwx**

**Shift Left Word [Immediate] [and Record]**

**e\_slwi** rA,rS,SH (Rc = 0)

**e\_slwi.** rA,rS,SH (Rc = 1)



```

if 'e_slwi[.]' then n ← SH
if se_slw then n ← GPR(RY)58:63
if se_slwi then n ← UI5
    
```



```

r ← ROTL32(GPR(RS or RX)32:63,n)
if n<32 then m ← MASK(32,63-n)
else          m ← 320
result32:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
    
```

Let the shift count *n* be the value specified by the contents of bits 58–63 of GPR(*rB* or *rY*), or by the value of the SH or UI5 field.

The contents of bits 32–63 of GPR(*rS* or *rX*) are shifted left *n* bits. Bits shifted out of position 32 are lost. Zeros are supplied to the vacated positions on the right. The 32-bit result is placed into bits 32–63 of GPR(*rA* or *rX*).


Shift amounts from 32 to 63 give a zero result.

Special Registers Altered: CR0 (if Rc = 1)



**System Call**

**se\_sc**



```

SRR1 ← MSR
SRR0 ← CIA+2
NIA ← IVPR32:47 || IVOR8,48:59 || 0b0000
MSRWE,EE,PR,IS,DS,FP,FE0,FE1 ← 0b0000_0000
    
```

**se\_sc** is used to request a system service. A system call interrupt is generated. The contents of the MSR are copied into SRR1 and the address of the instruction after the **se\_sc** instruction is placed into SRR0.

MSR[WE,EE,PR,IS,DS,FP,FE0,FE1] are cleared.

The interrupt causes the next instruction to be fetched from the address

```
IVPR[32–47]||IVOR8[48–59]||0b0000
```

This instruction is context synchronizing.

Special Registers Altered: SRR0 SRR1 MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]

**sraw**

Book E	User
--------	------

**sraw**

Shift right algebraic word [immediate]

**sraw** rA,rS,rB (Rc=0)  
**sraw.** rA,rS,rB (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1					rS			rA			rB					1	1	0	0	0	1	1	0	0	0	Rc

**srawi** rA,rS,SH (Rc=0)  
**srawi.** rA,rS,SH (Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1					rS			rA			SH					1	1	0	0	1	1	1	0	0	0	Rc

```

if 'sraw[.]' then n ← rB59:63
else n ← SH
r ← ROTL64(rS[32:63],64-n)
if 'sraw[.]' & rB58=1 then m ← 64'0
else m ← MASK(n+32,63)
s ← rS32
result0:63 ← r&m | (64's)&¬m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO

rA ← result0:63
CA ← s & ((r&¬m)32:63≠0)
    
```

If **sraw**[,], the shift count n is the contents of rB[58–63].

If **srawi**[,], the shift count n is the value of the SH field.

The contents of rS[32–63] are shifted right n bits. Bits shifted out of position 63 are lost. Bit 32 of rS is replicated to fill the vacated positions on the left. The 32-bit result is placed into rA[32–63]. rS[32] is replicated to fill bits rA[0–31].

CA is set if rS[32–63] contain a negative value and any 1 bits are shifted out of bit position 63; otherwise CA is cleared.

A shift amount of zero causes rA to receive EXTS(rS[32–63]), and CA to be cleared. For **sraw**[,] shift amounts from 32 to 63 give a result of 64 signed bits, and cause CA to receive rS[32] (that is, sign bit of rS[32–63]).

Other registers altered: CA CR0 (if Rc=1)





**\_srawx**

VLE	User
-----	------

**\_srawx**

**Shift Right Algebraic Word [Immediate] [and Record]**



```

if 'se_sraw' then n ← GPR(RY)59:63
if 'se_srawi' then n ← UI5
r ← ROTL32(GPR(RS or RX)32:63, 32-n)
if ((se_sraw & GPR(RY)58=1) then m ← 320
else m ← MASK(n+32, 63)
s ← GPR(RS or RX)32
result0:63 ← r&m | (32s)&¬m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
CA ← s & ((r&¬m)32:63≠0)
    
```

If **se\_sraw**, let the shift count *n* be the contents of bits 58–63 of GPR(*rY*).

If **se\_srawi**, let the shift count *n* be the value of the UI5 field.

The contents of bits 32–63 of GPR(*rS* or *rX*) are shifted right *n* bits. Bits shifted out of position 63 are lost. Bit 32 of *rS* or *rX* is replicated to fill vacated positions on the left. The 32-bit result is placed into bits 32–63 of GPR(*rA* or *rX*).

CA is set if bits 32–63 of GPR(*rS* or *rX*) contain a negative value and any 1 bits are shifted out of bit position 63; otherwise CA is cleared.

A shift amount of zero causes GPR(*rA* or *rX*) to receive EXTS(GPR(*rS* or *rX*)<sub>32:63</sub>), and CA to be cleared. For **se\_sraw**, shift amounts from 32 to 63 give a result of 64 sign bits, and cause CA to receive bit 32 of the contents of GPR(*rS* or *rX*) (that is, sign bit of GPR(*rS* or *rX*)<sub>32:63</sub>).

Special Registers Altered: CA  
 CR0 (if Rc = 1)




**srw** Book E User **srw**

**Shift right word**

**srw** rA,rS,rB (Rc=0)

**srw.** rA,rS,rB (Rc=1)



```

n ← rB59-63
r ← ROTL64(rS32-63,64-n)
if rB58=0 then m ← MASK(n+32,63)
else m ← 640
result0:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
rA ← result0:63
    
```

The shift count n is the value specified by the contents of rB[58–63].

The contents of rS[32–63] are shifted right n bits. Bits shifted out of position 63 are lost. Zeros are supplied to the vacated positions on the left. The 32-bit result is placed into rA[32–63]. Bits rA[0–31] are cleared.

Shift amounts from 32 to 63 give a zero result.




Other registers altered: CR0 (if Rc=1)

**\_srwx** VLE User **\_srwx**

**Shift Right Word [Immediate] [and Record]**

**e\_srwi** rA,rS,SH (Rc = 0)

**e\_srwi.** rA,rS,SH (Rc = 1)



```

n ← GPR(RB)59:63
if 'e_srwi[.]' then n ← SH
if 'se_srw' then n ← GPR(RY)59:63
if 'se_srwi' then n ← UI5
r ← ROTL32(GPR(RS or RX)32:63,32-n)
    
```

```

if ((se_srw & GPR(RY)58=1) then m ← 320
else      m ← MASK(n+32,63)
result32:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
    
```

If **e\_srw**, let the shift count *n* be the value of the SH field.

If **se\_srw**, let the shift count *n* be the contents of bits 58–63 of GPR(*rY*).

If **se\_srw**, let the shift count *n* be the value of the UI5 field.

The contents of bits 32–63 of GPR(*rS* or *rX*) are shifted right *n* bits. Bits shifted out of position 63 are lost. Zeros are supplied to the vacated positions on the left. The 32-bit result is placed into bits 32–63 of GPR(*rA* or *rX*).


Shift amounts from 32 to 63 give a zero result.

Special Registers Altered: CR0 (if Rc = 1)


**stb** Book E User **stb**

**Store byte [with update] [indexed]**

**stb** *rS*,*D*(*rA*) (D-mode, I=0)  
**stbu** *rS*,*D*(*rA*) (D-mode, I=1)



**stbx** *rS*,*rA*,*rB* (X-mode, U=0)  
**stbux** *rS*,*rA*,*rB* (X-mode, U=1)



```

if rA=0 then a ← 640 else a ← rA
if D-mode then EA ← 320 || (a + EXTS(D))32:63
if X-mode then EA ← 320 || (a + rB)32:63
MEM(EA,1) ← rS56:63
if U=1 then rA ← EA
    
```

The EA is calculated as follows:

- For **stb** and **stbu**, EA is bits 32–63 of the sum of the contents of *rA*, or 64 zeros if *rA*=0, and the sign-extended value of the D field.
- For **stbx** and **stbux**, EA is bits 32–63 of the sum of the contents of *rA*, or 64 zeros if *rA*=0, and the contents of *rB*.

The contents of rS[56–63] are stored into the byte addressed by EA.

If U=1 (with update), EA is placed into rA.


If U=1 (with update) and rA=0, the instruction form is invalid.

Other registers altered: None


**\_stbx** VLE User **\_stbx**

**Store Byte [with Update] [Indexed]**


**e\_stb** rS,D(rA) (D-mode)



**se\_stb** rZ,SD4(rX) (SD4-mode)



**e\_stbu** rS,D8(rA) (D8-mode)



if (RA=0 & lse\_stb) then a ← 320 else a ← GPR(RA or RX)

if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>

if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>

if SD4-mode then EA ← (a + (2<sup>8</sup>0 || SD4))<sub>32:63</sub>

MEM(EA, 1) ← GPR(RS or RZ)<sub>56:63</sub>

if e\_stbu then GPR(RA) ← EA

Let the EA be calculated as follows:

- For **e\_stb** and **e\_stbu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_stb**, let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field.

The contents of bits 56–63 of GPR(rS) are stored into the byte in memory addressed by EA.


- If **e\_stbu**, EA is placed into GPR(rA).
- If **e\_stbu** and rA = 0, the instruction form is invalid.
- None

**stfd** Book E User **stfd**


**Store floating-point double [with update] [indexed]**

**stfd** frS,D(rA) (D-mode, U=0)

**stfdu** frS,D(rA) (D-mode, U=1)



**stfdx** frS,rA,rB (X-mode, U=0)  
**stfdux** frS,rA,rB (X-mode, U=1)



if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))32:63  
 if X-mode then EA ← 320 || (a + rB)32:63  
 MEM(EA,8) ← frS  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **stfd** and **stfdu**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D instruction field.
- For **stfdx** and **stfdux**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The contents of frS are stored into the double word addressed by EA.


If U=1 (with update), EA is placed into rA.

If U=1 (with update) and rA=0, the instruction form is invalid.

If MSR[FP]=0, **stfd[u][x]** causes a floating-point unavailable interrupt.

Other registers altered: None

**stfiwx** Book E User **stfiwx**  
**Store floating-point as integer word indexed**  
**stfiwx** frS,rA,rB



if rA=0 then a ← 640 else a ← rA  
 EA ← 320 || (a + rB)32:63  
 MEM(EA,4) ← frS[32:63]

The EA is calculated as follows:

- For **stfiwx**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The contents of frS[32–63] are stored, without conversion, into the word addressed by EA.

If the contents of frS were produced, either directly or indirectly, by a load floating-point single instruction, a single-precision arithmetic instruction, or **frsp**, the value stored is

undefined. (The contents of **frS** are produced directly by such an instruction if **frS** is the target register for the instruction. The contents of **frS** are produced indirectly by such an instruction if **frS** is the final target register of a sequence of one or more floating-point move instructions, with the input to the sequence having been produced directly by such an instruction.)


If MSR[FP]=0, an attempt to execute **stfiwx** causes a floating-point unavailable interrupt.

Other registers altered: None


**stfs** Book E User **stfs**

**Store floating-point single [with update] [indexed]**

**stfs** frS,D(rA) (D-mode, U=0)  
**stfsu** frS,D(rA) (D-mode, U=1)



**stfsx** frS,rA,rB (X-mode, U=0)  
**stfsux** frS,rA,rB (X-mode, U=1)



if rA=0 then a ← 640 else a ← rA  
 if D-mode then EA ← 320 || (a + EXTS(D))32:63  
 if X-mode then EA ← 320 || (a + rB)32:63  
 MEM(EA,4) ← SINGLE(frS)  
 if U=1 then rA ← EA

The EA is calculated as follows:

- For **stfs** and **stfsu**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **stfsx** and **stfsux**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The contents of **frS** are converted to single format and stored into the word addressed by EA.

If U=1 (with update), EA is placed into rA.

If U=1 (with update) and rA=0, the instruction form is invalid.


If MSR[FP]=0, **stfs[u][x]** causes a floating-point unavailable interrupt.

Other registers altered: None


**sth** Book E User **sth**

**Store half word [with update] [indexed]**

**sth**  $rS, D(rA)$  (D-mode, U=0)  
**sth**  $rS, D(rA)$  (D-mode, U=1)



**sthx**  $rS, rA, rB$  (X-mode, U=0)  
**sthux**  $rS, rA, rB$  (X-mode, U=1)



if  $rA=0$  then  $a \leftarrow 640$  else  $a \leftarrow rA$   
 if D-mode then  $EA \leftarrow 320 \parallel (a + EXT(S,D))_{32:63}$   
 if X-mode then  $EA \leftarrow 320 \parallel (a + rB)_{32:63}$   
 $MEM(EA, 2) \leftarrow rS_{48:63}$   
 if  $U=1$  then  $rA \leftarrow EA$

The EA is calculated as follows:

- For **sth** and **sth**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D field.
- For **sthx** and **sthux**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

The contents of rS[48–63] are stored into the half word addressed by EA.

If U=1 (with update), EA is placed into rA.


If U=1 (with update) and rA=0, the instruction form is invalid.

Other registers altered: None


**\_sthx** VLE User **\_sthx**

**Store Halfword [with Update] [Indexed]**


**e\_sth**  $rS, D(rA)$  (D-mode)



**se\_sth**  $rZ, SD4(rX)$  (SD4-mode)



**e\_sthu**  $rS, D8(rA)$  (D8-mode)



if  $(RA=0 \ \& \ !se\_sth)$  then  $a \leftarrow 320$  else  $a \leftarrow GPR(RA \text{ or } RX)$

if D-mode then  $EA \leftarrow (a + EXTS(D))_{32:63}$   
 if D8-mode then  $EA \leftarrow (a + EXTS(D8))_{32:63}$   
 if SD4-mode then  $EA \leftarrow (a + (2^7 0 \parallel SD4 \parallel 0))_{32:63}$   
 $MEM(EA,2) \leftarrow GPR(RS \text{ or } RZ)_{48:63}$   
 if e\_sthu then  $GPR(RA) \leftarrow EA$

Let the EA be calculated as follows:

- For **e\_sth** and **e\_sthu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_sth** let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 1 bit.

The contents of bits 48–63 of GPR(rS) are stored into the half word in memory addressed by EA.

If **e\_sthu**, EA is placed into GPR(rA).

If **e\_sthu** and rA = 0, the instruction form is invalid.

Special Registers Altered: None


**sthbrx**

Book E	User
--------	------

**sthbrx**

**Store half word byte-reverse**

**sthbrx** rS,rA,rB



if rA=0 then  $a \leftarrow 64 0$  else  $a \leftarrow rA$   
 $EA \leftarrow 32 0 \parallel (a + rB)_{32:63}$   
 $MEM(EA,2) \leftarrow rS_{56:63} \parallel rS_{48:55}$

The EA is calculated as follows:

- For **sthbrx**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

rS[56–63] are stored into bits 0–7 of the half word addressed by EA. Bits 48–55 of rS are stored into bits 8–15 of the half word addressed by EA.

Other registers altered: None

Programming note: When EA references big-endian memory, these instructions have the effect of storing data in little-endian byte order. Likewise, when EA references little-endian memory, these instructions have the effect of storing data in big-endian byte order.





**stmw**

Book E	User
--------	------

**stmw**

**Store multiple word**

**stmw** **rS,D(rA)**

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31



```

if rA=0 then EA ← 320 || EXT(S,D)32:63
else EA ← 320 || (rA+EXT(S,D))32:63
r ← rS
do while r ≤ 31
    MEM(EA,4) ← GPR(r)32:63
    r ← r + 1
    EA ← 320 || (EA+4)32:63
    
```

The EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the sign-extended value of the D instruction field.

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined.

Other registers altered: None

**\_stmw**


VLE	User
-----	------

**\_stmw**

**Store Multiple Word**

**e\_stmw** **rS,D8(rA)** (D8-mode)

0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31



```

if RA=0 then EA ← EXT(S,D8)32:63
else EA ← (GPR(RA)+EXT(S,D8))32:63
r ← RS
do while r ≤ 31
    MEM(EA,4) ← GPR(r)32:63
    r ← r + 1
    EA ← (EA+4)32:63
    
```

Let the EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D8 instruction field.

Let n = (32 - rS). Bits 32–63 of registers GPR(rS) through GPR(31) are stored in n consecutive words in memory starting at address EA.

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined.

Special Registers Altered: None




**stswi**

Book E	User
--------	------


**stswi**

**Store string word (immediate | indexed)**

**stswi**                      rS,rA,NB



**stswx**                      rS,rA,rB



```

if rA=0 then a ← 640 else a ← rA
if 'stswi' then EA ← 320 || a32:63
if 'stswx' then EA ← 320 || (a + rB)32:63
if 'stswi' & NB=0 then n ← 32
if 'stswi' & NB≠0 then n ← NB
if 'stswx' then n ← XER57:63
r ← rS - 1
i ← 32
do while n > 0
    if i=32 then r ← r + 1 (mod 32)
    MEM(EA,1) ← GPR(r)i:i+7
    i ← i + 8
    if i = 64 then i ← 32
    EA ← 320 || (EA+1)32:63
    n ← n - 1
    
```

The EA is calculated as follows:

- For **stswi**, EA is 32 zeros concatenated with bits 32–63 of the contents of rA, or 32 zeros if rA=0.
- For **stswx**, EA is 32 zeros concatenated with bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

If **stswi**, let n=NB if NB≠0, n=32 if NB=0. If **stswx**, let n=XER[57–63]. n is the number of bytes to store. Let nr=CEIL(n÷4): nr is the number of registers to supply data.

n consecutive bytes starting at EA are stored from registers rS through GPR(rS+nr-1). Data is stored from the low-order 4 bytes of each GPR.

Bytes are stored left to right from each GPR. The register sequence can wrap to GPR0.

If **stswx** and n=0, no bytes are stored.

Other registers altered: None


Programming note: Store string word and load string word instructions allow movement of data between memory and registers without concern for alignment. They can be used for a short move between arbitrary locations or long moves between misaligned memory fields.




**stw** Book E User **stw**

**Store word [with update] [indexed]**

**stw**  $rS, D(rA)$  (D-mode, U=0)  
**stwu**  $rS, D(rA)$  (D-mode, U=1)



**stwx**  $rS, rA, rB$  (X-mode, U=0)  
**stwux**  $rS, rA, rB$  (X-mode, U=1)



if  $rA=0$  then  $a \leftarrow 64_0$  else  $a \leftarrow rA$   
 if D-mode then  $EA \leftarrow 32_0 \parallel (a + EXT(S,D))_{32:63}$   
 if X-mode then  $EA \leftarrow 32_0 \parallel (a + rB)_{32:63}$   
 $MEM(EA, 4) \leftarrow rS_{[32:63]}$   
 if  $U=1$  then  $rA \leftarrow EA$

The EA is calculated as follows:

- For **stw** and **stwu**, EA is bits 32–63 of the sum of the contents of  $rA$ , or 64 zeros if  $rA=0$ , and the sign-extended value of the D field.
- For **stwx** and **stwux**, EA is bits 32–63 of the sum of the contents of  $rA$ , or 64 zeros if  $rA=0$ , and the contents of  $rB$ .

The contents of  $rS[32–63]$  are stored into the word addressed by EA.

If  $U=1$  (with update), EA is placed into  $rA$ .


If  $U=1$  (with update) and  $rA=0$ , the instruction form is invalid.

Other registers altered: None


**\_stwx** VLE User **\_stwx**

**Store Word [with Update] [Indexed]**


**e\_stw**  $rS, D(rA)$  (D-mode)



**se\_stw**  $rZ, SD4(rX)$  (SD4-mode)



**e\_stwu**  $rS, D8(rA)$  (D8-mode)



if  $(RA=0 \ \& \ !se\_stw)$  then  $a \leftarrow 32_0$  else  $a \leftarrow GPR(RA \ \text{or} \ RX)$

if D-mode then  $EA \leftarrow (a + EXTS(D))_{32:63}$   
 if D8-mode then  $EA \leftarrow (a + EXTS(D8))_{32:63}$   
 if SD4-mode then  $EA \leftarrow (a + (2^6 0 \parallel SD4 \parallel 2^0))_{32:63}$   
 $MEM(EA,4) \leftarrow GPR(RS \text{ or } RZ)_{32:63}$

Let the EA be calculated as follows:

- For **e\_stw** and **e\_stwu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_stw**, let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 2 bits.

The contents of bits 32–63 of GPR(rS) are stored into the word in memory addressed by EA.

If **e\_stwu**, EA is placed into GPR(rA).

If **e\_stwu** and rA = 0, the instruction form is invalid.

Special Registers Altered: None


**stwbrx**

Book E	User
--------	------

**stwbrx**

**Store word byte-reverse**

**stwbrx**                                    rS,rA,rB



if rA=0 then  $a \leftarrow 64 0$  else  $a \leftarrow rA$   
 $EA \leftarrow 32 0 \parallel (a + rB)_{32:63}$   
 $MEM(EA,4) \leftarrow rS_{56:63} \parallel rS_{48:55} \parallel rS_{40:47} \parallel rS_{32:39}$

The EA is calculated as follows:

- For **stwbrx**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

Bits 56–63 of rS are stored into bits 0–7 of the word addressed by EA. Bits 48–55 of rS are stored into bits 8–15 of the word addressed by EA. Bits 40–47 of rS are stored into bits 16–23 of the word addressed by EA. Bits 32–39 of rS are stored into bits 24–31 of the word addressed by EA.

Other registers altered: None

Programming note: When EA references big-endian memory, these instructions have the effect of storing data in little-endian byte order. Likewise, when EA references little-endian memory, these instructions have the effect of storing data in big-endian byte order.


**stwcx.**

Book E	User
--------	------

**stwcx.**

**Store word conditional indexed**

**stwcx.** **rS,rA,rB**



```

if rA=0 then a ← 640 else a ← rA
EA ← 320 || (a + rB)32:63
if RESERVE then
    if RESERVE_ADDR = real_addr(EA) then
        MEM(EA,4) ← rS[32:63]
        CR0 ← 0b00 || 0b1 || XERSO
    else
        u ← undefined 1-bit value
        if u then MEM(EA,4) ← rS[32:63]
        CR0 ← 0b00 || u || XERSO
    RESERVE ← 0
else
    CR0 ← 0b00 || 0b0 || XERSO
    
```

The EA is calculated as follows:

- For **stwcx.**, EA is bits 32–63 of the sum of the contents of rA, or 64 zeros if rA=0, and the contents of rB.

If a reservation exists and the address specified by the **stwcx.** is the same as that specified by the **lwarx** instruction that established the reservation, the contents of rS[32–63] are stored into the word addressed by EA and the reservation is cleared.

If a reservation exists but the address specified by **stwcx.** is not the same as that specified by the load and reserve instruction that established the reservation, the reservation is cleared, and it is undefined whether the instruction completes without altering memory.

If a reservation does not exist, the instruction completes without altering memory.

CR field 0 is set to reflect whether the store operation was performed, as follows:

$$CR0[LT,GT,EQ,SO] = 0b00 || store\_performed || XER[SO]$$

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined.

Other registers altered: CR0

Programming notes:

- **stwcx.**, in combination with **lwarx**, permits the programmer to write a sequence of instructions that appear to perform an atomic update operation on a memory location. This operation depends on a single reservation resource in each processor. At most one reservation exists on any given processor: there are not separate reservations for words and for double words.
- Because **stwcx.** instructions have implementation dependencies (such as the granularity at which reservations are managed), they must be used with care. The operating system should provide system library programs that use these instructions to implement the high-level synchronization functions (such as, test and set, and compare



and swap) needed by application programs. Application programs should use these library programs, rather than use **stwcx.** directly.

- The granularity with which reservations are managed is implementation-dependent. Therefore, the memory to be accessed by **stwcx.** should be allocated by a system library program. Additional information can be found in [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#)
- When correctly used, the load and reserve and store conditional instructions can provide an atomic update function for a single aligned word (**lwarx** and **stwcx.**) of memory. In general, correct use requires that **lwarx** be paired with **stwcx.** with the same address specified by both instructions of the pair. The only exception is that an unpaired **stwcx.** to any (scratch) effective address can be used to clear any reservation held by the processor. Examples of correct uses of these instructions to emulate primitives such as fetch and add, test and set, and compare and swap can be found in [Appendix C: Programming examples](#).

A reservation is cleared if any of the following events occur:

- The processor holding the reservation executes another load and reserve instruction; this clears the first reservation and establishes a new one.
- The processor holding the reservation executes a store conditional instruction to any address.
- Another processor executes any store instruction to the address associated with the reservation.
- Any mechanism, other than the processor holding the reservation, stores to the address associated with the reservation.

See [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#), for additional information.


**\_sub**

VLE	User
-----	------

**\_sub**

**Subtract**

**se\_sub**



$$\text{sum}_{32:63} \leftarrow \text{GPR}(\text{RX}) + \neg\text{GPR}(\text{RY}) + 1$$

$$\text{GPR}(\text{RX}) \leftarrow \text{sum}_{32:63}$$

The sum of the contents of GPR(rX), the one's complement of contents of GPR(rY), and 1 is placed into GPR(rX).

Special Registers Altered: None


**subf**

Book E	User
--------	------

**subf**

**Subtract from**

<b>subf</b>	rD,rA,rB	(OE=0, Rc=0)
<b>subf.</b>	rD,rA,rB	(OE=0, Rc=1)
<b>subfo</b>	rD,rA,rB	(OE=1, Rc=0)
<b>subfo.</b>	rD,rA,rB	(OE=1, Rc=1)



```

carry0:63 ← Carry(¬rA + rB + 1)
sum0:63 ← ¬rA + rB + 1
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
rD ← sum
    
```

The sum of the one's complement of the contents of rA, the contents of rB, and 1 is placed into rD.

Other registers altered:

- CR0 (if Rc=1)
- SO OV (if OE=1)


**\_subfx**

VLE	User
-----	------

**\_subfx**

**Subtract From**

**se\_subf** rX,rY



```

sum32:63 ← ¬GPR(RX) + GPR(RY) + 1
GPR(RX) ← sum32:63
    
```

The sum of the one's complement of the contents of GPR(rX), the contents of GPR(rY), and 1 is placed into GPR(rX).

Special Registers Altered: None

**subfc**

Book E	User
--------	------

**subfc**

**Subtract from carrying**

<b>subfc</b>	rD,rA,rB	(OE=0, Rc=0)
<b>subfc.</b>	rD,rA,rB	(OE=0, Rc=1)
<b>subfco</b>	rD,rA,rB	(OE=1, Rc=0)
<b>subfco.</b>	rD,rA,rB	(OE=1, Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	1	1	1	1	1					rD			rA								rB			OE	0	0	0	0	0	1	0	0	0	Rc

```

carry0:63 ← Carry(¬rA + rB + 1)
sum0:63 ← ¬rA + rB + 1
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)

if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO

rD ← sum
CA ← carry32
    
```

The sum of the one's complement of the contents of rA, the contents of rB, and 1 is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)

**subfe**

Book E	User
--------	------

**subfe**

**Subtract from extended**

<b>subfe</b>	rD,rA,rB	(OE=0, Rc=0)
<b>subfe.</b>	rD,rA,rB	(OE=0, Rc=1)
<b>subfeo</b>	rD,rA,rB	(OE=1, Rc=0)
<b>subfeo.</b>	rD,rA,rB	(OE=1, Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	1	1	1	1	1					rD			rA									rB			OE	0	1	0	0	0	1	0	0	0	Rc

```

if E=0 then Cin ← CA
carry0:63 ← Carry(¬rA + rB + Cin)
sum0:63 ← ¬rA + rB + Cin
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)

if Rc=1 then do
    LT ← sum32:63 < 0
    
```





GT  $\leftarrow$  sum<sub>32:63</sub> > 0  
 EQ  $\leftarrow$  sum<sub>32:63</sub> = 0  
 CR0  $\leftarrow$  LT || GT || EQ || SO

rD  $\leftarrow$  sum  
 CA  $\leftarrow$  carry<sub>32</sub>

For **subfe[o][.].**, the sum of the one's complement of the contents of rA, the contents of rB, and CA is placed into rD.


Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)

**subfic** Book E | User **subfic**

**Subtract from immediate carrying**

**subfic** rD,rA,SIMM



carry<sub>0:63</sub>  $\leftarrow$  Carry( $\neg$ rA + EXTS(SIMM) + 1)  
 sum<sub>0:63</sub>  $\leftarrow$   $\neg$ rA + EXTS(SIMM) + 1  
 rD  $\leftarrow$  sum  
 CA  $\leftarrow$  carry<sub>32</sub>


The sum of the one's complement of the contents of rA, the sign-extended value of the SIMM field, and 1 is placed into rD.

Other registers altered: CA

**\_subficx** VLE | User **\_subficx**

**Subtract From Immediate Carrying [and Record]**

**e\_subfic** rD,rA,SCI8 (Rc = 0)  
**e\_subfic.** rD,rA,SCI8 (Rc = 1)



imm  $\leftarrow$  SCI8(F,SCL,UI8)  
 carry<sub>32:63</sub>  $\leftarrow$  Carry( $\neg$ GPR(RA) + imm + 1)  
 sum<sub>32:63</sub>  $\leftarrow$   $\neg$ GPR(RA) + imm + 1  
 if Rc=1 then do  
     LT  $\leftarrow$  sum<sub>32:63</sub> < 0  
     GT  $\leftarrow$  sum<sub>32:63</sub> > 0  
     EQ  $\leftarrow$  sum<sub>32:63</sub> = 0  
     CR0  $\leftarrow$  LT || GT || EQ || SO

GPR(RD) ← sum<sub>32:63</sub>  
 CA ← carry<sub>32</sub>

The sum of the one's complement of the contents of GPR(rA), the value of SCI8, and 1 is placed into GPR(rD).

Special Registers Altered: CA CR0 (if Rc=1)

**subfme**

Book E	User
--------	------

**subfme**

**Subtract from minus one extended**

<b>subfme</b>	rD,rA	(OE=0, Rc=0)
<b>subfme.</b>	rD,rA	(OE=0, Rc=1)
<b>subfmeo</b>	rD,rA	(OE=1, Rc=0)
<b>subfmeo.</b>	rD,rA	(OE=1, Rc=1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31		
0	1	1	1	1	1					rD						rA					///		OE	0	1	1	1	0	1	0	0	0	Rc

if E=0 then Cin ← CA  
 carry<sub>0:63</sub> ← Carry(¬rA + Cin + 0xFFFF\_FFFF\_FFFF\_FFFF)  
 sum<sub>0:63</sub> ← ¬rA + Cin + 0xFFFF\_FFFF\_FFFF\_FFFF  
 if OE=1 then do  
     OV ← carry<sub>32</sub> ⊕ carry<sub>33</sub>  
     SO ← SO | (carry<sub>32</sub> ⊕ carry<sub>33</sub>)  
 if Rc=1 then do  
     LT ← sum<sub>32:63</sub> < 0  
     GT ← sum<sub>32:63</sub> > 0  
     EQ ← sum<sub>32:63</sub> = 0  
     CR0 ← LT || GT || EQ || SO  
 rD ← sum  
 CA ← carry<sub>32</sub>

For **subfme[o][.]**, the sum of CA, <sup>64</sup>1, and the one's complement of the contents of rA is placed into rD.

Other registers altered:

- CA  
     CR0 (if Rc=1)  
     SO OV (if OE=1)


**subfze**

Book E	User
--------	------

**subfze**

**Subtract from zero extended**

<b>subfze</b>	rD,rA	(OE=0, Rc=0)
<b>subfze.</b>	rD,rA	(OE=0, Rc=1)
<b>subfzeo</b>	rD,rA	(OE=1, Rc=0)
<b>subfzeo.</b>	rD,rA	(OE=1, Rc=1)



```

if E=0 then Cin ← CA
carry0:63 ← Carry(−rA + Cin)
sum0:63 ← −rA + Cin
if OE=1 then do
    OV ← carry32 ⊕ carry33
    SO ← SO | (carry32 ⊕ carry33)

if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO

rD ← sum
CA ← carry32
    
```

For **subfze[o][.]**, the sum of the one’s complement of the contents of rA and CA is placed into rD.

Other registers altered:

- CA
- CR0 (if Rc=1)
- SO OV (if OE=1)


**\_subix**

VLE	User
-----	------

**\_subix**

**Subtract Immediate [and Record]**

<b>se_subi</b>	rX,OIMM	(Rc = 0)
<b>se_subi.</b>	rX,OIMM	(Rc = 1)



1. OIMM = OIM5 + 1

```

sum32:63 ← GPR(RX) + −(270 || OFFSET(OIM5)) + 1
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    
```



$$CR0 \leftarrow LT \parallel GT \parallel EQ \parallel SO$$

$$GPR(RX) \leftarrow \text{sum}_{32:63}$$

The sum of the contents of GPR(rX), the one's complement of the zero-extended value of the offseted OIM5 field (a final value in the range 1–32), and 1 is placed into GPR(rX).

Special Registers Altered: CR0 (if Rc = 1)

**tlbivax**

Book E	Supervisor
--------	------------

**tlbivax**

**TLB Invalidate virtual address indexed**

**tlbivax**

rA,rB

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	///					rA				rB				1	1	0	0	0	1	0	0	1	0	/		

$$\begin{aligned} &\text{if } rA=0 \text{ then } a \leftarrow {}^{64}0 \text{ else } a \leftarrow rA \\ &EA \leftarrow {}^{32}0 \parallel (a + rB)_{32:63} \\ &AS \leftarrow \text{implementation-dependent value} \\ &\text{ProcessID} \leftarrow \text{implementation-dependent value} \\ &VA \leftarrow AS \parallel \text{ProcessID} \parallel EA \\ &\text{InvalidateTLB}(VA) \end{aligned}$$

EIS note: Executing **tlbivax** invalidates any TLB entry that corresponds to a virtual address calculated by this instruction if IPROT is not set; this includes invalidating TLB entries on other devices as well as on the processor executing **tlbivax**. Thus an invalidate operation is broadcast throughout the coherent domain of the processor executing **tlbivax**. On some implementations, HID1[ABE] must be set to allow management of external L2 caches (for implementations with L2 caches) as well as other L1 caches in the system.

EA calculation: for rA≠0	Addressing ModeEA for rA=0EA
$\parallel (rA+rB)_{32:63}$	${}^{32}0 \parallel rB_{32:63} {}^{32}0$

Address space (AS) is defined as implementation-dependent (for example, it could be MSR[DS] or a bit from an implementation-dependent SPR).

ProcessID is implementation-dependent (for example, it could be from the PID or from an implementation-dependent SPR). The EIS implements the architected PID and additional implementation-specific PIDs. See [Section 3.12.1: Process ID registers \(PID0–PIDn\)](#).

The virtual address (VA) is the value AS || ProcessID || EA.

A TLB entry corresponding to VA is made invalid (that is, removed from the TLB). This instruction causes the target TLB entry to be invalidated in all processors.


The operation performed by this instruction is ordered by **mbar** (or **msync**) with respect to a subsequent **tlbsync** executed by the processor executing **tlbivax**. Operations caused by **tlbivax** and **tlbsync** are ordered by **mbar** as a set of operations independent of the other sets that **mbar** orders.

Other registers altered: None

Programming notes:

- The effects of the invalidation are not guaranteed to be visible to the programming model until the completion of a context synchronizing operation. See [Section 4.2.3.6: Context synchronization](#).
- Care must be taken not to invalidate TLB entries that contain interrupt vector mappings.

**tlbre** Book E Supervisor **tlbre**  
**TLB Read entry**  
**tlbre**



1. This field is defined as allocated by the Book E architecture, for possible use in an implementation. These bits are not implemented by the EIS.

The RTL for the EIS definition of **tlbre** is as follows:

```

tlb_entry_id = MAS0(TLBSEL, ESEL | MAS2(EPN)
result = MMU(tlb_entry_id)
MAS0, MAS1, MAS2, MAS3, (and MAS7 if HID0[EN_MAS7_UPDATE] = 1) = result
    
```

Bits 6–20 of the encoding are allocated for implementation-dependent use and may be used to specify the source TLB entry, the source portion of the source TLB entry, and the target resource into which the result is placed. The EIS makes no use of these bits.

The implementation-defined TLB entry is read, and the implementation-defined portion of the TLB entry is extracted and placed into an implementation-defined target resource.

If the instruction specifies a TLB entry that does not exist, the results are undefined.


EIS implementation note: **tlbre** causes the contents of a single TLB entry to be extracted from the MMU and be placed in the corresponding fields of the MMU assist (MAS) registers. The entry extracted is specified by the TLBSEL, ESEL and EPN fields of MAS0 and MAS2. The contents extracted from the MMU are placed in MAS0–MAS3.

See the user’s manual for the implementation.

Execution of this instruction is restricted to supervisor mode.

Other registers altered: MAS0, MAS1, MAS2, and MAS3, as defined by the EIS

**tlbsx** Book E Supervisor **tlbsx**  
**TLB Search indexed**  
**tlbsx** **rA,rB**



1. This field is defined as allocated by the Book E architecture, for possible use in an implementation. These bits are not implemented by the EIS.

if RA!=0 then generate exception  
 EA =  $32_0 \parallel \text{GPR}(\text{RB})_{32:63}$   
 ProcessID = MAS6(SPID)  
 AS = MAS6(SAS)  
 VA0 = AS  $\parallel$   $(\text{MMUCFG}[\text{PIDSIZE}] + 1)_0 \parallel \text{EA}$   
 VA1 = AS  $\parallel$  ProcessID  $\parallel$  EA  
 if Valid\_TLB\_matching\_entry\_exists (VA0) or Valid\_TLB\_matching\_entry\_exists (VA1)  
 #  
 #  
 MAS0, MAS1, MAS2, MAS3 = result

EA calculation:      Addressing ModeEA for rA=0EA for rA≠0  
 $32_0 \parallel \text{rB}_{32:63} \parallel 32_0 \parallel (\text{rA} + \text{rB})_{32:63}$

Note that rA = 0 is a preferred form for **tlbsx** and that some ST implementations take an illegal instruction exception program interrupt if rA != 0.

Virtual address 0 (VA0) is the value AS  $\parallel$   $(\text{MMUCFG}[\text{PIDSIZE}] + 1)_0 \parallel \text{EA}$   
 Virtual address 1 (VA1) is the value AS  $\parallel$  ProcessID  $\parallel$  EA

If the TLB contains an entry corresponding to VA, an implementation-dependent value is placed into an implementation-dependent-specified target. Otherwise the contents of the implementation-dependent-specified target are left undefined.

Other registers altered: implementation-dependent. See [Section 4.3.2.4: Supervisor-level tlb management instructions](#).


**tlbsync**

Book E	Supervisor
--------	------------

**tlbsync**

**TLB Synchronize**

**tlbsync**



**tlbsync** provides an ordering function for the effects of all **tlbivax** instructions executed by the processor executing **tlbsync**, with respect to the memory barrier created by a subsequent **msync** instruction executed by the same processor. Executing **tlbsync** ensures that all of the following occur.

- All TLB invalidations caused by **tlbivax** instructions preceding the **tlbsync** instruction will have completed on any other processor before any memory accesses associated with data accesses caused by instructions following the **msync** instruction are performed with respect to that processor.
- All memory accesses by other processors for which the address was translated using the translations being invalidated, will have been performed with respect to the processor executing the **msync** instruction, to the extent required by the associated memory-coherence required attributes, before the **mbar** or **msync** instruction's memory barrier is created.

The operation performed by this instruction is ordered by the **mbar** and **msync** instructions with respect to preceding **tlbivax** instructions executed by the processor executing the



**tlbsync** instruction. The operations caused by **tlbivax** and **tlbsync** are ordered by **mbar** as a set of operations that is independent of the other sets that **mbar** orders.

The **tlbsync** instruction may complete before operations caused by **tlbivax** instructions preceding the **tlbsync** instruction have been performed.


Execution of this instruction is restricted to supervisor mode.

Other registers altered: None

**tlbwe**

Book E	Supervisor
--------	------------

**tlbwe**  
**TLB Write entry**  
**tlbwe**



1. This field is defined as allocated by the Book E architecture, for possible use in an implementation. These bits are not implemented by the EIS.

Bits 6–20 of the instruction encoding are allocated for implementation-dependent use, and may be used to specify the target TLB entry, the target portion of the target TLB entry, and the source of the value that is to be written into the TLB. The EIS does not make use of these bits.

The contents of the implementation-dependent–specified source are written into the implementation-dependent–specified portion of the implementation-dependent–specified TLB entry.

If the instruction specifies a TLB entry that does not exist, the results are undefined.

Execution of this instruction may cause other implementation-dependent effects. See the user’s manual for the implementation.

Execution of this instruction is restricted to supervisor mode.

Other registers altered: None


Programming notes:

- The effects of the update are not guaranteed to be visible to the programming model until the completion of a context synchronizing operation. See [Section 4.2.3.6: Context synchronization](#).
- Care must be taken not to invalidate any TLB entry that contains the mapping for any interrupt vector.


**tw** Book E User **tw**

**Trap word [immediate]**

**tw** TO,rA,rB



**twi** TO,rA,SIMM



```

a ← EXTS(rA32:63)
if 'tw' then b ← EXTS(rB32:63)
if 'twi' then b ← EXTS(SIMM)
if (a < b) & TO0 then TRAP
if (a > b) & TO1 then TRAP
if (a = b) & TO2 then TRAP
if (a <u b) & TO3 then TRAP
if (a >u b) & TO4 then TRAP
    
```

For **tw**, the contents of rA[32–63] are compared with the contents of rB[32–63].

For **twi**, the contents of rA[32–63] are compared with the sign-extended value of the SIMM field.


If any bit in the TO field is set and its corresponding condition is met by the result of the comparison, then the system trap handler is invoked.

Other registers altered: None


**wrtee** Book E Supervisor **wrtee**

**Write MSR external enable [immediate]**

**wrtee** rS



**wrteei** E



```

if 'wrtee' then MSR[EE] ← rS48
if 'wrteei' then MSR[EE] ← E
    
```

For **wrtee**, rS[48] is placed into MSR[EE].

For **wrteei**, the value specified in the E field is placed into MSR[EE].



Execution of this instruction is restricted to supervisor mode.

In addition, changes to MSR[EE] are effective as soon as the instruction completes. Thus if MSR[EE]=0 and an external interrupt is pending, executing a **wrttee** or **wrtteei** that sets MSR[EE] causes the external interrupt to be taken before the next instruction is executed, if no higher priority exception exists.

Other registers altered: MSR

Programming note: **wrttee** and **wrtteei** are used to update of MSR[EE] without affecting other MSR bits. Typical usage is as follows:


```

mfmsr          Rn          #save EE in GPR(Rn)48
wrtteei        0           #turn off EE
:              :           :
:              :           :
:              :           :
wrttee         Rn          #restore EE without altering other MSR bits
that may have changed
    
```


**xor** Book E | User **xor**

**XOR [Immediate [shifted]]**

**xor** rA,rS,rB (Rc=0)  
**xor.** rA,rS,rB (Rc=1)



**xori** rA,rS,UIMM (S=0, Rc=0)  
**xoris** rA,rS,UIMM (S=1, Rc=0)



```

if 'xori' then b ← 480 || UIMM
if 'xoris' then b ← 320 || UIMM || 160
if 'xor[.]' then b ← rB
result0:63 ← rS ⊕ b
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
rA ← result
    
```

For **xori**, the contents of rS are XORed with <sup>48</sup>0 || UIMM.  
 For **xoris**, the contents of rS are XORed with <sup>32</sup>0 || UIMM || <sup>16</sup>0.  
 For **xor[.]**, the contents of rS are XORed with the contents of rB.  
 The result is placed into rA.

Other registers altered: CR0 (if Rc=1)

**\_xorx**


VLE	User
-----	------

**\_xorx**

**XOR [Immediate] [and Record]**

**e\_xori** **rA,rS,SCI8** (**Rc = 0**)

**e\_xori.** **rA,rS,SCI8** (**Rc = 1**)



if 'e\_xori[.]' then  $b \leftarrow \text{SCI8}(F,SCL,UI8)$

$\text{result}_{32:63} \leftarrow \text{GPR}(RS) \oplus b$

if Rc=1 then do

$LT \leftarrow \text{result}_{32:63} < 0$

$GT \leftarrow \text{result}_{32:63} > 0$

$EQ \leftarrow \text{result}_{32:63} = 0$

$CR0 \leftarrow LT \parallel GT \parallel EQ \parallel SO$

$\text{GPR}(RA) \leftarrow \text{result}$

For **e\_xori[.]**, the contents of GPR(rS) are XORed with SCI8.

The result is placed into GPR(rA).

Special Registers Altered: CR0 (if Rc = 1)

## 8 Auxiliary processing units (APUs)

This chapter describes the APUs defined by the EIS, which are as follows:

- [Section 8.1: Integer select APU](#)
- [Section 8.2: Performance monitor APU](#)
- [Section 8.3: Signal processing engine APU \(SPE APU\)](#)
- [Section 8.4: Embedded vector and scalar single-precision floating-point APUs \(SPFP APUs\)](#)
- [Section 8.5: Machine check APU](#)
- [Section 8.6: Debug APU](#)
- [Section 8.7: Alternate time base](#)

Note that individual processors may implement APUs that are not defined by the EIS. Individual processors may either further extend these APUs or may implement a subset of the resources described here. See the documentation for the individual implementation.

### 8.1 Integer select APU

Control code, which is characterized by unpredictable short branches, is common in embedded applications. When mispredicted, these branches cause long pipeline delays. The integer select (**isel**) APU consists of a single instruction (**isel**), a conditional register move that helps eliminate some of these branches. The **isel** instruction works as follows:

```

if crB then
    rD = rA
else
    rD = rB

```

The **isel** instruction allows more efficient implementation of a condition sequence such as the one in the following generic example:

```

int16 global1,..., global37,...;
....
void procedure17(int16 parm) {
    if      (global1 == 27) {
        global37 = parm + 17;
    }
    else {
        global37 = parm - 17;
    }
}

```

#### 8.1.1 Integer select APU programming model

The integer select APU includes only the **isel** instruction, described in [Chapter 7: Instruction set on page 334](#). It accesses the GPRs and the CR and does not implement additional registers or interrupt resources.

### 8.1.2 Using isel to Improve conditional branch performance

The Integer Select instruction, **isel**, can be used to handle short conditional branch segments more efficiently. **isel** has two source registers and one destination register. Under the control of a specified condition code bit, it copies one or the other source operand to the destination.

[Table 208](#) shows a coding example with and without the **isel** instruction.

**Table 208. Recoding with isel**

Code sequence without isel	Code sequence with isel
<pre> cmpi cr3, r17, 27; bne cr3, NotEqual; addi r15, r17, 17; jmp Assign; NotEqual:     addi r15, r17, -17; Assign:     stw r15, (rGlobals + g37);                     </pre>	<pre> cmpi cr3, r17, 27; addi r15, r17, 17; addi r16, r17, -17; isel r15, r15, r16, cr3.eq stw r15, (rGlobals + g37);                     </pre>

The sequence without **isel** turns conditional branches into a code sequence that sets a condition code according to the results of a computation. It uses a conditional branch to choose a target sequence, but needs an unconditional branch for the IF clause. The conditional branch is often hard to predict, the code sequences are generally small, and the resulting throughput is typically low.

The sequence using **isel** does the following:

- Sets a condition code according to the results of a comparison
- Has code that executes both the IF and the ELSE segments
- Has a final statement that copies the results of one of the segments to the desired destination register
- Works well for small code segments and for unpredictable branches
- Can reduce code size

## 8.2 Performance monitor APU

The EIS defines the performance monitor as an APU. Software communication with the performance monitor APU is achieved through performance monitor registers (PMRs) rather than SPRs. The PMRs are used for enabling conditions that can trigger an APU-defined performance monitor interrupt.

### 8.2.1 Performance monitor APU programming model

The performance monitor APU provides a set of PMRs for defining, enabling, and counting conditions that trigger the performance interrupt. The APU defines instructions for reading and writing the PMRs.

### 8.2.1.1 Performance monitor APU registers

The performance monitor APU defines IVOR35 (SPR 531) for indicating the address of the performance monitor interrupt vector. IVOR35 is described in [Interrupt vector offset registers \(IVORs\) on page 85](#).

The APU also defines a set of PMRs that are separate from the SPR resources. However, like SPRs and as shown in [Table 209](#) and [Table 210](#), bit 5 indicates whether a register is user- or supervisor-accessible. Supervisor-level PMRs in [Table 209](#) are accessed through the **mtpmr** and **mfpmr** instructions. Attempting to read or write supervisor-level registers while in user-mode causes a privilege exception.

**Table 209. Performance monitor registers—supervisor level**

Register name	Abbreviation	PMR number	pmr[0–4]	pmr[5–9]
Counter 0	PMC0	16	00000	10000
Counter 1	PMC1	17	00000	10001
Counter 2	PMC2	18	00000	10010
Counter 3	PMC3	19	00000	10011
Local control a0	PMLCa0	144	00100	10000
Local control a1	PMLCa1	145	00100	10001
Local control a2	PMLCa2	146	00100	10010
Local control a3	PMLCa3	147	00100	10011
Local control b0	PMLCb0	272	01000	10000
Local control b1	PMLCb1	273	01000	10001
Local control b2	PMLCb2	274	01000	10010
Local control b3	PMLCb3	275	01000	10011
Global control 0	PMGC0	400	01100	10000

The user-level PMRs in [Table 210](#) are read-only and are accessed with the **mfpmr** instruction. Attempting to write user-level registers in either supervisor or user mode causes an illegal instruction exception.

**Table 210. Performance monitor registers—user level (read-only)**

Register name	Abbreviation	PMR number	pmr[0–4]	pmr[5–9]
Counter 0	UPMC0	0	00000	00000
Counter 1	UPMC1	1	00000	00001
Counter 2	UPMC2	2	00000	00010
Counter 3	UPMC3	3	00000	00011
Local control a0	UPMLCa0	128	00100	00000
Local control a1	UPMLCa1	129	00100	00001
Local control a2	UPMLCa2	130	00100	00010
Local control a3	UPMLCa3	131	00100	00011

**Table 210. Performance monitor registers—user level (read-only) (continued)**

Register name	Abbreviation	PMR number	pmr[0–4]	pmr[5–9]
Local control b0	UPMLCb0	256	01000	00000
Local control b1	UPMLCb1	257	01000	00001
Local control b2	UPMLCb2	258	01000	00010
Local control b3	UPMLCb3	259	01000	00011
Global control 0	UPMGC0	384	01100	00000

PMRs are fully described in [Section 3.16: Performance monitor registers \(PMRs\)](#).

### 8.2.1.2 Performance monitor apu instructions

The APU also defines the instructions in [Table 211](#) to move to and move from these PMRs. Full descriptions of these instructions can be found in [Chapter 7 on page 334](#).

**Table 211. Performance monitor apu instructions**

Name	Mnemonic	Syntax
Move from Performance Monitor Register	<b>mfpmr</b>	rD,PMRN
Move to Performance Monitor Register	<b>mtpmr</b>	PMRN,rS

### 8.2.1.3 Performance monitor APU interrupt model

The performance monitor APU provides a performance monitor interrupt that is triggered by an enabled condition or event.

## 8.3 Signal processing engine APU (SPE APU)

This section describes the SPE APU programming model, exceptions, and functions.

### 8.3.1 Overview

This section describes the instruction set architecture of the signal processing engine (SPE) APU. The SPE APU 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. SPE also architects an accumulator register to allow for back-to-back operations without loop unrolling.

### 8.3.2 Nomenclature and conventions

Several conventions regarding nomenclature are used in this document:

- The signal processing engine APU is abbreviated as SPE.
- All register bit numbering is 64-bit, with bit 0 being the most significant bit. Registers that are only 32-bit define bit 32 as the most significant bit. For both 32- and 64-bit registers, bit 63 is the least significant bit.
- Bits 0 to 31 of a 64-bit register are referenced as upper word, even word or high word element of the register. Bits 32–63 are referred to as lower word, odd word, or low word element of the register. Each half is an element of a 64-bit GPR.
- 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 referenced as odd half words.
- Mnemonics for SPE instructions generally begin with the letters ‘ev’ (embedded vector).

### 8.3.3 Programming model

This section describes SPE registers, instructions, and interrupts.


#### 8.3.3.1 General operation

SPE instructions generally take elements from each source register and operate on them with the corresponding elements of a second source register (and/or the accumulator) to produce results. Results are placed in the destination register and/or the accumulator. Instructions that are vector in nature (that is, they produce results of more than one element) provide results for each element that are independent of the computation of the other elements. These instructions can also be used to perform scalar DSP operations by ignoring the results of the upper 32-bit half of the register file.

There are no record forms of SPE instructions. SPE compare instructions store the compare result into the condition register (CR). The meaning of the CR bits is now overloaded for SPE operations. SPE compare instructions specify a CR field, two source registers, and the type of compare: greater than, less than, or, equal. Two bits of the CR field are written with the result of the vector compare, one for each element. The remaining two bits reflect the ANDing and ORing of the vector compare results.


#### 8.3.3.2 GPR registers

The SPE APU requires a GPR register file with thirty-two 64-bit registers. For 32-bit implementations, PowerPC Book E 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 the SPE APU was not implemented. SPE APU instructions view the 64-bit register as being composed of a vector of two elements, each of which is 32 bits wide. (Some instructions read or write 16-bit elements.) The most significant 32 bits are called the upper word, high word or even word. The least significant 32 bits are called the lower word, low word or odd word. Unless otherwise specified, SPE instructions write all 64 bits of the destination register.



**8.3.3.3 Accumulator register**

A partially visible accumulator register (ACC) is provided for the integer/fractional multiply accumulate (MAC) forms of instructions. The accumulator is a 64-bit register that holds the results of the multiply accumulate forms of SPE fixed-point instructions. The accumulator 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. The accumulator is partially visible to the programmer in the sense that its results do not have to be explicitly read to use them. Instead they are always copied into a 64-bit destination GPR, which is specified as part of the instruction. Based upon the type of instruction, the accumulator can hold either a single 64-bit value or a vector of two 32-bit elements.



**8.3.3.4 Signal processing embedded floating-point status and control register (SPEFSCR)**

Status and control for SPE uses the SPEFSCR, described in [Chapter 3.14.1: Signal processing, embedded floating-point status, control register \(SPEFSCR\) on page 120](#). The embedded floating-point APUs also use SPEFSCR. Status and control bits are shared for embedded floating-point operations and SPE vector operations. The SPEFSCR is implemented as SPR number 512 and is read and written by the **mfspr** and **mtspr** instructions in both user and supervisor mode.

**8.3.3.5 SPE exception bit in ESR**

ESR[SPE] is defined as the SPE exception bit. This bit is set whenever the processor takes an interrupt related to the execution of SPE instructions. (Note that the same bit is used for embedded floating-point APU exceptions. Thus, SPE and embedded floating-point exceptions are indistinguishable in the ESR.)

**8.3.3.6 SPE available bit in MSR**

MSR[SPE] is defined as the SPE available bit. If this bit is not set and software attempts to execute an SPE instruction, the SPE APU unavailable interrupt is taken.

Software note: This bit can be used by software 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.

**8.3.3.7 Data formats**

The SPE APU provides two different data formats, integer and fractional. Both data formats can be treated as signed or unsigned quantities.

**8.3.3.8 Integer format**

Integer data format is the same as what is conventionally used in computing.

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.

**8.3.3.9 Fractional format**

Fractional data format is the same that is conventionally used for DSP fractional arithmetic. Fractional data is useful for representing data converted from analog devices.

Unsigned fractions consist of 16-, 32-, or 64-bit binary fractional values that range from 0 to less than 1. Unsigned fractions place the decimal point immediately to the left of the most significant bit. The most significant bit of the value represents the value  $2^{-1}$ , the next most significant bit 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 set OV or OVH in SPEFSCR. SPE does not contain explicit instructions that manipulate unsigned fractional data. Unsigned integer forms produce the same bit results as unsigned fractional values would; therefore, unsigned fractional instruction forms are not defined for SPE.


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 most significant bit. 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  $-1$  set OV or OVH in the SPEFSCR. 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 least-significant bit (lsb) of the shifted result.

**8.3.3.10 Computational operations**

SPE supports several different computational capabilities. These can be grouped 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 242* shows how operations are typically performed in vector operations.

**Figure 242. Two-element vector operations**



- Multiply and accumulate instructions. These instructions perform multiply operations, add the result to the accumulator and place the result into the destination register and the accumulator. These instructions are composed of different multiply forms, data

formats, and data accumulate options. The mnemonics for these instructions indicate their various characteristics. These are shown in [Table 212](#).

**Table 212. Mnemonic extensions for multiply accumulate instructions**

Extension	Meaning	Comments
<b>Multiply form</b>		
<b>he</b>	Half word even	16 X 16 → 32
<b>heg</b>	Half word even guarded	16 X 16 → 32, 64-bit final accum result
<b>ho</b>	Half word odd	16 X 16 → 32
<b>hog</b>	Half word odd guarded	16 X 16 → 32, 64-bit final accum result
<b>w</b>	Word	32 X 32 → 64
<b>wh</b>	Word high	32 X 32 → 32 (high order 32 bits of product)
<b>wl</b>	Word low	32 X 32 → 32 (low order 32 bits of product)
<b>Data format</b>		
<b>smf</b>	Signed modulo fractional	Modulo, no saturation or overflow
<b>smi</b>	Signed modulo integer	Modulo, no saturation or overflow
<b>ssf</b>	Signed saturate fractional	Saturation on product and accumulate
<b>ssi</b>	Signed saturate integer	Saturation on product and accumulate
<b>umi</b>	Unsigned modulo integer	Modulo, no saturation or overflow
<b>usi</b>	Unsigned saturate integer	Saturation on product and accumulate
<b>Accumulate option</b>		
<b>a</b>	Place in accumulator	Result → accumulator
<b>aa</b>	Add to accumulator	Accumulator + result → accumulator
<b>aaw</b>	Add to accumulator	Accumulator <sub>0:31</sub> + result <sub>0:31</sub> → accumulator <sub>0:31</sub> Accumulator <sub>32:63</sub> + result <sub>32:63</sub> → accumulator <sub>32:63</sub>
<b>an</b>	Add negated to accumulator	Accumulator – result → accumulator
<b>anw</b>	Add negated to accumulator	Accumulator <sub>0:31</sub> – result <sub>0:31</sub> → accumulator <sub>0:31</sub> Accumulator <sub>32:63</sub> – result <sub>32:63</sub> → accumulator <sub>32:63</sub>

- 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.

### 8.3.3.11 SPE exceptions and interrupts

The APU defines the following SPE exceptions:

- SPE/embedded floating-point unavailable exception (causes the SPE/embedded floating point unavailable interrupt)
- SPE vector alignment exception (causes the alignment interrupt)

Interrupt vector offset registers (IVORs) IVOR32 (SPE/embedded floating-point unavailable interrupt) and IVOR5 (alignment interrupt) are used by the interrupt model. The SPR number for IVOR32 is 528; IVOR5 is defined by Book E. These registers are privileged.

SPE/Embedded floating point unavailable exception

The SPE/embedded floating point unavailable exception occurs when execution of an SPE instruction (except **brinc**) is attempted and bit 38 (SPE available, MSR[SPE]) is not set. If the SPE/embedded floating point unavailable exception occurs, a SPE/embedded floating point unavailable exception interrupt is taken and the processor suppresses execution of the instruction causing the exception. SRR0, SRR1, MSR, and ESR 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[36] bit is set. All other ESR bits are cleared.

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

Software note: This exception is also used by the embedded floating-point APUs in the same manner. It should be used by software 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.

SPE vector alignment exception

The SPE vector alignment exception is taken if the EA of any of the following instructions is not aligned to a 64-bit boundary: **evladd**, **evlddx**, **evldw**, **evldwx**, **evldh**, **evldhx**, **evstdd**, **evstdx**, **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[56] bit 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 used in the load or the store.

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

### 8.3.3.12 Interrupt priorities

The following list shows the priority order in which SPE APU and SPFP APU interrupts are taken (see [Embedded floating-point interrupts on page 681](#)):

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

### 8.3.4 Instruction definitions

[Chapter 7: Instruction set on page 334](#), gives complete descriptions of SPE and embedded floating-point instructions. [Section 7.3.1: SPE APU 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.

## 8.4 Embedded vector and scalar single-precision floating-point APUs (SPFP APUs)

This section describes the instruction set architecture of the embedded floating-point APUs. The EIS defines the following APUs:

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

Each of these APUs may be implemented independently of the other. In addition, there is a strong relationship with the SPE APU in that each of the embedded floating-point APUs shares a common status register with the SPE.

### 8.4.1 Nomenclature and conventions

Several conventions regarding nomenclature are used in this document:

- The embedded vector single-precision floating-point APU operations are abbreviated as vector floating-point or vector SPFP.
- The embedded scalar single-precision floating-point APU operations are abbreviated as scalar SPFP.
- The embedded scalar double-precision floating-point APU operations are abbreviated as scalar DPFP.
- Bits 0 to 31 of a 64-bit register are referenced as field 0, upper half, upper word, or high-word element of the register. Bits 32–63 are referred to as field 1, lower half, or lower-word element of the register. Each half is an element of a 64-bit GPR.
- Mnemonics for vector floating-point instructions generally begin with the letters ‘**evf**’ (embedded vector float).
- Mnemonics for single-precision floating-point instructions generally begin with the letters ‘**efs**’ (embedded floating single).
- References to ‘floating-point’ or ‘embedded SPFP’ refer to both APUs.

### 8.4.2 Embedded floating-point APUs programming model

The embedded floating-point APUs use the GPRs as source and destination operands; however, double precision and vector instruction require 64-bit GPRs as described in [Embedded floating-point APUs GPR implementations on page 680](#).

### 8.4.2.1 Embedded floating-point instructions

The following sections show opcodes for the three embedded floating-point APUs, as follows:

- [Opcodes for embedded vector floating-point instructions on page 677](#)
- [Opcodes for embedded scalar single-precision floating-point instructions on page 677](#)
- [Opcodes for embedded scalar double-precision floating-point instructions on page 678](#)

Opcodes for embedded vector floating-point instructions

[Table 213](#) lists the embedded vector floating-point opcodes.

**Table 213. Embedded vector floating-point instruction opcodes**

Instruction	Opcode bits					Comments
	0–5	6–10	11–15	16–20	21–31	
<b>evfsabs</b>	4	rD	rA	00000	010 1000 0100	
<b>evfsadd</b>	4	rD	rA	rB	010 1000 0000	
<b>evfscfsf</b>	4	rD	00000	rB	010 1001 0011	
<b>evfscfsi</b>	4	rD	00000	rB	010 1001 0001	
<b>evfscfuf</b>	4	rD	00000	rB	010 1001 0010	
<b>evfscfui</b>	4	rD	00000	rB	010 1001 0000	
<b>evfscmpeq</b>	4	crfD 00	rA	rB	010 1000 1110	
<b>evfscmpgt</b>	4	crfD 00	rA	rB	010 1000 1100	
<b>evfscmplt</b>	4	crfD 00	rA	rB	010 1000 1101	
<b>evfsctsf</b>	4	rD	00000	rB	010 1001 0111	
<b>evfsctsi</b>	4	rD	00000	rB	010 1001 0101	
<b>evfsctsiz</b>	4	rD	00000	rB	010 1001 1010	
<b>evfsctuf</b>	4	rD	00000	rB	010 1001 0110	
<b>evfsctui</b>	4	rD	00000	rB	010 1001 0100	
<b>evfsctuiz</b>	4	rD	00000	rB	010 1001 1000	
<b>evfsdiv</b>	4	rD	rA	rB	010 1000 1001	
<b>evfsmul</b>	4	rD	rA	rB	010 1000 1000	
<b>evfsnabs</b>	4	rD	rA	00000	010 1000 0101	
<b>evfsneg</b>	4	rD	rA	00000	010 1000 0110	
<b>evfssub</b>	4	rD	rA	rB	010 1000 0001	rA - rB
<b>evfststeq</b>	4	crfD 00	rA	rB	010 1001 1110	
<b>evfststgt</b>	4	crfD 00	rA	rB	010 1001 1100	
<b>evfststlt</b>	4	crfD 00	rA	rB	010 1001 1101	

Opcodes for embedded scalar single-precision floating-point instructions

[Table 214](#) lists the embedded scalar single-precision floating-point opcodes.

**Table 214. Embedded scalar single-precision floating-point instruction opcodes**

Instruction	Opcode bits					Comments
	0–5	6–10	11–15	16–20	21–31	
<b>efsabs</b>	4	rD	rA	00000	010 1100 0100	
<b>efsadd</b>	4	rD	rA	rB	010 1100 0000	
<b>efscfd</b>	4	rD	00000	rB	010 1100 1111	
<b>efscfsf</b>	4	rD	00000	rB	010 1101 0011	
<b>efscfsi</b>	4	rD	00000	rB	010 1101 0001	
<b>efscfuf</b>	4	rD	00000	rB	010 1101 0010	
<b>efscfui</b>	4	rD	00000	rB	010 1101 0000	
<b>efscmpeq</b>	4	crfD 00	rA	rB	010 1100 1110	
<b>efscmpgt</b>	4	crfD 00	rA	rB	010 1100 1100	
<b>efscmplt</b>	4	crfD 00	rA	rB	010 1100 1101	
<b>efscfsf</b>	4	rD	00000	rB	010 1101 0111	
<b>efscfsi</b>	4	rD	00000	rB	010 1101 0101	
<b>efscfsiz</b>	4	rD	00000	rB	010 1101 1010	
<b>efscfuf</b>	4	rD	00000	rB	010 1101 0110	
<b>efscfui</b>	4	rD	00000	rB	010 1101 0100	
<b>efscfui</b>	4	rD	00000	rB	010 1101 1000	
<b>efscfui</b>	4	rD	rA	rB	010 1100 1001	
<b>efsmul</b>	4	rD	rA	rB	010 1100 1000	
<b>efsnabs</b>	4	rD	rA	00000	010 1100 0101	
<b>efsneg</b>	4	rD	rA	00000	010 1100 0110	
<b>efssub</b>	4	rD	rA	rB	010 1100 0001	rA - rB
<b>efststeq</b>	4	crfD 00	rA	rB	010 1101 1110	
<b>efststgt</b>	4	crfD 00	rA	rB	010 1101 1100	
<b>efststlt</b>	4	crfD 00	rA	rB	010 1101 1101	

Opcodes for embedded scalar double-precision floating-point instructions

[Table 215](#) lists the embedded scalar double-precision floating-point opcodes.

**Table 215. Embedded scalar double-precision floating-point instruction opcodes**

Instruction	Opcode bits					Comments
	0–5	6–10	11–15	16–20	21–31	
<b>efdabs</b>	4	rD	rA	00000	010 1110 0100	
<b>efdadd</b>	4	rD	rA	rB	010 1110 0000	
<b>efdcfs</b>	4	rD	00000	rB	010 1110 1111	

Table 215. Embedded scalar double-precision floating-point instruction opcodes (continued)

Instruction	Opcode bits					Comments
	0–5	6–10	11–15	16–20	21–31	
<b>efdcfsf</b>	4	rD	00000	rB	010 1111 0011	
<b>efdcfsi</b>	4	rD	00000	rB	010 1111 0001	
<b>efdcfsid</b>	4	rD	00000	rB	010 1110 0011	64-bit only
<b>efdcfuf</b>	4	rD	00000	rB	010 1111 0010	
<b>efdcfui</b>	4	rD	00000	rB	010 1111 0000	
<b>efdcfuid</b>	4	rD	00000	rB	010 1110 0010	64-bit only
<b>efdcmpcq</b>	4	crfD 00	rA	rB	010 1110 1110	
<b>efdcmpgt</b>	4	crfD 00	rA	rB	010 1110 1100	
<b>efdcmplt</b>	4	crfD 00	rA	rB	010 1110 1101	
<b>efdctsf</b>	4	rD	00000	rB	010 1111 0111	
<b>efdctsi</b>	4	rD	00000	rB	010 1111 0101	
<b>efdctsidz</b>	4	rD	00000	rB	010 1110 1011	64-bit only
<b>efdctsiz</b>	4	rD	00000	rB	010 1111 1010	
<b>efdctuf</b>	4	rD	00000	rB	010 1111 0110	
<b>efdctui</b>	4	rD	00000	rB	010 1111 0100	
<b>efdctuidz</b>	4	rD	00000	rB	010 1110 1010	64-bit only
<b>efdctuiz</b>	4	rD	00000	rB	010 1111 1000	
<b>efddiv</b>	4	rD	rA	rB	010 1110 1001	
<b>efdmul</b>	4	rD	rA	rB	010 1110 1000	
<b>efdnabs</b>	4	rD	rA	00000	010 1110 0101	
<b>efdneg</b>	4	rD	rA	00000	010 1110 0110	
<b>efdsb</b>	4	rD	rA	rB	010 1110 0001	rA - rB
<b>efdsteq</b>	4	crfD 00	rA	rB	010 1111 1110	
<b>efdstgt</b>	4	crfD 00	rA	rB	010 1111 1100	
<b>efdstlt</b>	4	crfD 00	rA	rB	010 1111 1101	

## Optional load/store instructions

All embedded floating-point APUs use GPRs to hold and operate on floating-point values. The APUs do not architect load and store instructions to move the data to and from memory, but instead rely on existing instructions in the architecture to perform this function. In the case where either the vector single-precision embedded floating-point APU or the scalar double-precision embedded floating-point APU is implemented on a 32-bit implementation, the GPRs are required to be 64-bits long. Because a 32-bit implementation contains no load or store instructions that operate on 64-bit data, new instructions are required to perform these actions. In this case (and for a 64-bit implementation), an implementation may implement the following load/store instructions from the SPE APU.

For scalar double-precision:

- **evladd**—Vector load doubleword into doubleword
- **evlddx**—Vector load doubleword into doubleword indexed
- **evstd**—Vector store doubleword of doubleword
- **evstdx**—Vector store doubleword of doubleword
- **evmergehi**—Vector merge high
- **evmergelo**—Vector merge low

For vector single-precision, all of the vector load/store word and doubleword instructions, merge instructions, and word forms of splat instructions may be implemented. Because the vector single-precision embedded floating-point APU uses a significant set of the SPE vector load/store/merge instructions, it is strongly recommended that the SPE APU be present when implementing the vector single-precision embedded floating-point APU.

Floating-point conversion models

Each APU contains 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 [Section 7.3.2: Embedded floating-point conversion models](#), as a group of functions that is called from the individual instruction pseudo-RTL descriptions included in the instruction descriptions in [Chapter 7: Instruction set on page 334](#).

#### 8.4.2.2 Embedded floating-point registers

The embedded floating-point APUs share register resources with the SPE APU, as described in the following sections.

Embedded floating-point APUs GPR implementations

Embedded floating-point operations are performed in the GPRs of the processor.

The vector floating-point and double-precision floating-point require a GPR register file with thirty-two 64-bit registers. This is consistent with the SPE APU. Thus, these can coexist with the SPE APU.

Single-precision floating-point requires a GPR register 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.

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 as well, 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.



Signal processing embedded floating-point status and control register (SPEFSCR)

The embedded floating-point APUs use the SPEFSCR, which is described in [Section 3.14.1: Signal processing, embedded floating-point status, control register \(SPEFSCR\)](#). The SPE APU also uses SPEFSCR. Status and control bits are shared for vector floating-point operations, single-precision floating-point operations and SPE vector operations. The SPEFSCR is implemented as SPR number 512 and is read and written by **mf spr** and **mt spr** in both user and supervisor mode. Vector floating-point instructions affect both the high- and low-element floating-point status flags (bits 34–39 and 50–55). Scalar SPFP instructions affect only the low-element flags and leave the high element flags undefined.

Embedded floating-point exception bit—ESR[SPE]

ESR[SPE] is defined as the embedded floating-point exception bit. This bit is set whenever the processor takes an interrupt related to the execution of the embedded floating-point instructions. (Note that the same bit is used for SPE APU exceptions. Thus, SPE and embedded floating-point interrupts are indistinguishable in the ESR.)

### 8.4.2.3 Embedded floating-point interrupts

The following sections describe the embedded floating-point APU interrupts:

- [SPE/embedded floating-point unavailable interrupt on page 681](#)
- [Embedded floating-point data interrupt on page 681](#)
- [Embedded floating-point round interrupt on page 682](#)

SPE/embedded floating-point unavailable interrupt

The SPE/embedded floating-point unavailable interrupt vector is used by the embedded scalar double-precision floating-point APU and the embedded vector single-precision floating-point APU. It is not used by the embedded scalar single-precision floating-point APU. The SPE/embedded floating-point unavailable interrupt occurs when an embedded vector floating-point or an embedded scalar double-precision floating-point instruction is executed and bit 38 of the MSR is not set. If the SPE/embedded floating-point unavailable interrupt occurs, the processor suppresses execution of the instruction causing the exception.

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[24] is set. All other ESR bits are cleared.

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

This interrupt is also used by the SPE APU in the same manner. It should be used by software 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.

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[SPE] 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]||IVOR32[48–59]||0b0000.

#### 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.

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 might occur regardless of the value of FINXE unless an embedded floating-point data interrupt also occurs and is taken.

When the 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 both zero.

The FG and FX bits are provided so that an interrupt handler can round the result as it desires. FG (the guard bit) is the value of the bit immediately to the right of the least significant bit of the destination format mantissa from the infinitely precise intermediate calculation before rounding. FX (the sticky bit) is the value of the OR of all bits to the right of the guard bit (FG) 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[SPE] is set. All other ESR bits are cleared.
- SPEFSCR FGH, FG, FXH, and FX are set appropriately. SPEFSCR[FINXS] is set.

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

### 8.4.2.4 Interrupt priorities

The following list shows the priority order in which SPE and embedded floating-point interrupts are taken (see [Interrupt priorities on page 675](#)):

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

An embedded floating-point data interrupt is taken if either element of a vector or scalar floating-point operation generates an embedded floating-point data exception. An embedded floating-point round interrupt is taken if either element of a vector floating-point operation or a scalar floating-point operation generates an embedded floating-point round exception and no operation (both element for vector floating-point) generates an embedded floating-point data exception.

### 8.4.3 Embedded floating-point APU operations

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

#### 8.4.3.1 Operational modes

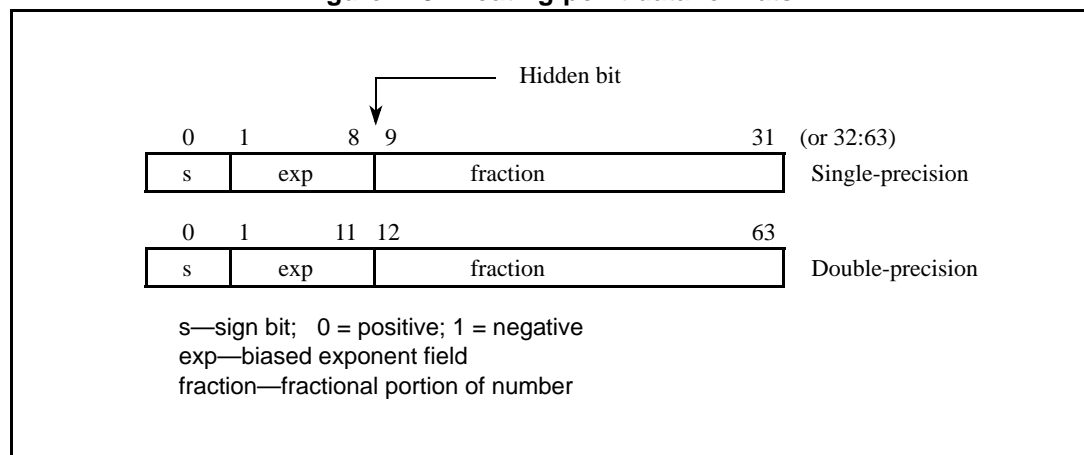
All embedded floating-point operations are governed by the setting of the mode bit in SPEFSCR. The mode bit defines how floating-point results are computed and how floating-point exceptions are handled. Mode 0 defines a real-time, default-results-oriented mode that saturates results. Other modes are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

#### 8.4.3.2 Floating-point data formats

Single-precision floating-point data elements are 32 bits wide with 1 sign bit (s), 8 bits of biased exponent (exp) and 23 bits of fraction.

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 243. Floating-point data formats



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  $+127$ . 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 ( $p_{\max}$ ) is represented by the encoding 0x7F7F\_FFFF, which is approximately  $3.4E+38$  ( $2^{128}$ ), and the minimum positive normalized value ( $p_{\min}$ ) is represented by the encoding 0x0080\_0000, which is approximately  $1.2E-38$  ( $2^{-126}$ ).

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

Zeros of both positive and negative sign are represented by a biased exponent value ( $e$ ) of zero and a fraction that is zero.

Infinities of both positive and negative sign are represented by a maximum exponent field value (255 for single-precision) and a fraction that is zero.

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

Not-a-Numbers (NaNs) are represented by a maximum exponent field value (255 for single-precision) and a fraction that is non-zero.

### 8.4.3.3 Overflow and underflow

Defining  $p_{\max}$  to be the most positive normalized value (farthest from zero),  $p_{\min}$  the smallest positive normalized value (closest to zero),  $n_{\max}$  the most negative normalized value (farthest from zero) and  $n_{\min}$  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  $r > p_{\max}$  or  $r < n_{\max}$ . 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  $A_{\text{exp}} \geq 254 \mid B_{\text{exp}} \geq 254$  then overflow

double-precision addition:

if  $A_{\text{exp}} \geq 2046 \mid B_{\text{exp}} \geq 2046$  then overflow

single-precision multiplication:

if  $A_{\text{exp}} + B_{\text{exp}} - 127 \geq 254$  then overflow

double-precision multiplication:

if  $A_{\text{exp}} + B_{\text{exp}} - 1023 \geq 2046$  then overflow

An underflow is said to have occurred if the numerically correct result of an instruction is such that  $0 < r < p_{\text{min}}$  or  $n_{\text{min}} < r < 0$ . 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  $A_{\text{exp}} - B_{\text{exp}} - 127 \leq 1$  then underflow

double-precision multiplication:

if  $A_{\text{exp}} - B_{\text{exp}} - 1023 \leq 1$  then underflow

The embedded floating-point APUs will not produce  $+\text{Inf}$ ,  $-\text{Inf}$ , NaN, or a Denormalized number. If the result of an instruction overflows and floating-point overflow exceptions are disabled (SPEFSCR[FOVFE] is cleared),  $p_{\text{max}}$  or  $n_{\text{max}}$  is generated as the result of that instruction depending upon 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.

#### 8.4.3.4 IEEE 754 compliance

The embedded floating-point APU implements a floating-point system as defined in ANSI/IEEE Standard 754-1985 but may rely on software support in order to conform fully with the standard. Thus, whenever an input operand of a floating-point instruction has data values that are  $+\text{infinity}$ ,  $-\text{infinity}$ , denorm, or NaN, or when the result of an operation produces an overflow or an underflow, an interrupt may be taken and the interrupt handler is responsible for delivering IEEE 754-compliant behavior if desired.

When floating-point invalid input exceptions are disabled (SPEFSCR[FINVE] is cleared), default results are provided by the hardware when an infinity, denorm, or NaN input is received, or for the operation  $0/0$ . When floating-point underflow exceptions are disabled (SPEFSCR[FUNFE] is cleared) and the result of a floating-point operation underflows, a signed zero result is produced. The inexact exception is also signaled for this condition. When floating-point overflow exceptions are disabled (EFSCR[FOVFE] is cleared) and the result of a floating-point operation overflows, a  $p_{\text{max}}$  or  $n_{\text{max}}$  result is produced. The inexact exception is also signaled for this condition. An exception enable flag (SPEFSCR[FINXE]) is also provided for generating an interrupt when an inexact result is produced, to allow a software handler to conform to the IEEE 754 standard. A divide-by-zero exception enable flag (SPEFSCR[FDBZE]) is provided for generating an interrupt when a divide-by-zero operation is attempted to allow a software handler to conform to the IEEE 754 standard. All of these exceptions may be disabled, and the hardware then delivers 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  $-\text{infinity}$ . Thus,  $-0 + -0 = -0$  is the only case in which the result is a  $-0$ ; all other cases that result in a zero value give  $+0$  unless the rounding mode is round to  $-\text{infinity}$ .

Note that when exceptions are disabled and default results computed, operations having input values that are denormalized may provide different bit-exact 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 on other implementations.

#### 8.4.3.5 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.

Not all sticky bits are required to be updated by an implementation. Only the FINXS and FDBZS sticky bits are required to be set by hardware. Thus for FINVS, FUNFS and FOVFS, software is required to perform sticky bit setting unless software knows that a given implementation updates them in hardware. This can be achieved by enabling the appropriate exceptions and performing the sticky bit updating in the software interrupt handler. If an implementation provides sticky bit handling for any sticky bits other than FINXS and FDBZS, it must provide it for all sticky bits.

#### 8.4.4 Implementation options summary

There are several options that may be chosen for a given implementation. This section summarizes all the items that are implementation dependent and should be used to help decide which implementation dependent features are chosen.

- APUs. Each of the APUs can be implemented independently of one another. The vector single-precision floating-point APU should be implemented only if the SPE APU is implemented; however, this is not required.
- Both the vector single-precision floating-point APU and the scalar double-precision floating-point APU allow the optional implementation of 64-bit load and store instructions as well as merge upper and lower instructions from the SPE APU. This allows data to be moved in and out of the upper half of a register for 32-bit implementations with 64-bit registers.
- Overflow and underflow conditions may be signaled by doing exponent evaluation of the operation. If by examining the exponents, an overflow or underflow could occur, the implementation may choose to signal an overflow or underflow. It is recommended that future implementations do not use this estimation and 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, negative absolute value operations, an implementation may choose to either simply perform the sign bit operation ignoring exceptions, or to compute the operation and handle exceptions and saturation where appropriate.
- The FGH and FXH bits of the SPEFSCR are undefined upon the 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 conversions to integer values.

## 8.5 Machine check APU

The machine check APU defines features for the machine check interrupt in addition to those defined by the PowerPC architecture and the Book E version of the PowerPC architecture. The machine check APU includes an enhanced definition of the machine check interrupt type similar to the Book E–defined critical interrupt.

### 8.5.1 Machine check APU programming model

The APU defines dedicated save and restore SPRs, MSRR0 and MSRR1, so a machine check interrupt does not affect the CSRR0, CSRR1, or ESR registers as defined by the Book E architecture.

The APU also defines a separate Return from Machine Check Interrupt instruction, **rfmci**, that restores context from MSRR0 and MSRR1 when the machine check interrupt handler completes.

#### 8.5.1.1 Machine check APU register model

The machine check APU defines different register for the machine check interrupt resources than the Book E definition. These are as follows:

- Machine-check save/restore register 0 (MCSRR0)—SPR 570. Holds the instruction where fetching begins after **rfmci** executes, typically at the end of the machine check interrupt handler. See [Machine check save/restore register 0 \(MCSRR0\) on page 89](#).
- Machine-check save/restore register 1 (MCSRR1)—SPR 571. Holds the machine state copied to the MSR when a machine check interrupt occurs. The MCSRR1 value is restored to the MSR when **rfmci** executes, typically at the end of the machine check interrupt handler. See [Machine check save/restore register 1 \(MCSRR1\) on page 90](#).
- Machine check syndrome register (MCSR)—SPR 572. MCSR has fields that identify causes for a machine check interrupt along with an indication of whether the processor can recover from the machine check interrupt. See [Machine check syndrome register \(MCSR\) on page 90](#).

Note, however, that the MSR[ME] bit, defined by the original PowerPC architecture, is also used in Book E and in the machine check APU to enable the machine check interrupt.

### 8.5.1.2 Machine check APU instruction model

The Return from Machine Check Interrupt instruction, **rfmci**, is context-synchronizing; it works its way to the final execute stage, updates architected registers, and redirects instruction flow. When **rfmci** executes, data is restored from MCSRR0 and MCSRR1. The **rfi** and **rftci** instructions do not affect MCSRR0 and MCSRR1. This instruction is described in [Chapter 4: Instruction model on page 134](#).

### 8.5.1.3 Machine check interrupt

The machine check APU is consistent with the machine check exception as defined in Book E with the following differences:

- Machine check is no longer a critical interrupt but uses MCSRR0 and MCSRR1 for saving the return address and the MSR in case the machine check is recoverable.
- The Return from Machine Check Interrupt instruction (**rfmci**) is implemented to support the return to the address saved in MCSRR0.
- The machine check syndrome register, MCSR, is used (instead of ESR) to log the cause of the machine check.

## 8.6 Debug APU

This section describes the instruction set architecture of software accessible debug related items for Book E Implementations (EIS).

The debug APU defines an additional interrupt class for debug interrupts. This allows the debug features to be used in the software that is providing service for critical class interrupts. This is accomplished by providing specific save and restore registers for debug interrupts and providing a new return from interrupt instruction (return from debug interrupt).

The debug APU reassigns debug interrupts into its own interrupt class, adding a new set of registers used to save the machine context upon the occurrence of a debug interrupt, and adds a new instruction, Return From Debug Interrupt (**rfdi**), to return from a debug interrupt and restore the machine state from the new set of registers. This APU redefines PowerPC Book E debug interrupt behavior.

An implementation may choose to provide the debug APU and also provide a method to disable the debug APU, reverting to using the critical interrupt as defined in Book E. If such a capability is provided, HIDO[DAPUEN] should be implemented.

### 8.6.1 Debug APU programming model

The following sections described the debug APU's extensions to the Book E interrupt, register, and interrupt models.

### 8.6.2 Debug APU register model

The debug interrupt defines the following registers:

- Debug save/restore register 0 (DSRR0). When a debug interrupt is taken, DSRR0 is set to the current or next instruction address. When **rfdi** is executed, instruction execution continues at the address in DSRR0.
- Debug save/restore register 1 (DSRR1). When a debug interrupt is taken, the contents of the MSR are placed into DSRR1. When **rfdi** is executed, the contents of DSRR1 are



placed into the MSR. Bits of DSRR1 that correspond to reserved bits in the MSR are also reserved.

This instruction is fully described in [Chapter 7: Instruction set on page 334](#).

The debug APU defines fields in the following Book E–defined registers:

- Debug status register (DBSR). New event fields, described in [Table 216](#), have been added to DBSR to record critical interrupt taken events and critical interrupt return events.

**Table 216. EIS-defined DBSR field descriptions**

Bits	Name	Description
57	CIRPT	Critical interrupt taken debug event. A critical interrupt taken debug event occurs when DBCR0[CIRPT] = 1 and a critical interrupt (any interrupt that uses the critical class, that is, uses CSRR0 and CSRR1) occurs. 0 No critical interrupt taken debug event has occurred. 1 A critical interrupt taken debug event occurred.
58	CRET	Critical interrupt return debug event. A critical interrupt return debug event occurs when DBCR0[CRET] = 1 and a return from critical interrupt (an <b>rfci</b> instruction is executed) occurs. 0 No critical interrupt return debug event has occurred. 1 A critical interrupt return debug event occurred.

- The debug control register 0 (DBCR0), The debug APU adds event enable bits to DBCR0, described in [Table 217](#), to control critical interrupt taken events, and critical interrupt return events.

**Table 217. DBCR0 field descriptions**

Bits	Name	Description
57	CIRPT	Critical interrupt taken debug event. A critical interrupt taken debug event occurs when DBCR0[CIRPT] = 1 and a critical interrupt (any interrupt that uses the critical class, that is, uses CSRR0 and CSRR1) occurs. 0 Critical interrupt taken debug events are disabled. 1 Critical interrupt taken debug events are enabled.
58	CRET	Critical interrupt return debug event. A critical interrupt return debug event occurs when DBCR0[CRET] = 1 and a return from critical interrupt (an <b>rfci</b> instruction is executed) occurs. 0 Critical interrupt return debug events are disabled. 1 Critical interrupt return debug events are enabled.

### 8.6.3 Debug APU instruction model

The debug APU defines the **supervisor-level rfdi** instruction to restore state after a debug interrupt. The contents of DSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address DSRR0[0–61]||0b00. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0, CSRR0, or DSRR0 by the interrupt processing mechanism is the address of the instruction that would have been

executed next had the interrupt not occurred (that is, the address in DSRR0 at the time of the execution of the **rfdi**). This instruction is fully described in [Chapter 7](#).

### 8.6.3.1 Debug APU interrupt model

A debug interrupt occurs when no higher priority exception exists, a debug exception is presented to the interrupt mechanism, and MSR[DE] = 1. The specific cause or causes of debug exceptions are unchanged from Book E.

DSRR0, DSRR1, MSR, debug address register, and debug status register are updated as follows:

Debug save/restore register 0 (DSRR0) is set to an instruction address. DSRR0 is set to the EA of an instruction that was executing or just completed execution when the debug exception occurred. DSRR0 is set the same as CSRR0 is defined to be set in Book E on a debug interrupt. CSRR0 is not changed as the result of a debug interrupt.

Debug save/restore register 1 (DSRR1) is set to the contents of the MSR at the time of the interrupt. CSRR1 is not changed as the result of a debug interrupt.

MSR[CM] is set to the value of MSR[ICM]. MSR[ICM] and MSR[ME] are unchanged and all other defined MSR bits are cleared.

The DBSR and the debug control registers (DBCR0–DBCR2) operate as described in Book E with the addition of a critical interrupt taken debug event and a critical return debug event.

Instruction execution resumes at address IVPR[0–47]||IVOR15[48–59]||0b0000.

## 8.7 Alternate time base

The alternate time base APU defines a time base counter similar to the time base defined in the PowerPC architecture. It is intended to be used for measuring time in implementation defined intervals. It differs from the time base defined by the PowerPC architecture in that it is not writable and always counts up, wrapping when the 64-bit count overflows.

### 8.7.1 Programming model

The alternate time base is simply a 64-bit counter that counts up at some implementation dependent rate. Although not required, it is recommended that the rate be at the core clock frequency or as small a multiple of the frequency as practical by the implementation. Consult the user documentation for devices that support this feature.

The counter can be read by executing an **mf spr** instruction specifying the ATB (or ATBL) register, but cannot be written. In 32-bit mode, reading the ATB (or ATBL) register will place the lower 32 bits of the counter into the target register. In 64-bit mode all 64 bits of the counter are placed in the target register. A second SPR register ATBU, is defined that accesses only the upper 32 bits of the counter. Thus the upper 32 bits of the counter may be read into a register by reading the ATBU register regardless of computation mode.

The alternate time base is analogous to the time base in the PowerPC architecture except that it counts at a different frequency and is not writable.

The effect of power savings mode or core frequency changes on counting in the alternate time base is implementation dependent. See the user document for details.

Implementation Note: An implementation may choose to directly alias the alternate time base to the time base counter if the granularity of time base counting is acceptable.

### 8.7.1.1 Registers

The programming model consists of two SPRs, alternate time base lower and upper (ATBL and ATBU).

Alternate time base registers (ATBL and ATBU)

The ATBL and ATBU registers are described in [Chapter 3.15: Alternate time base registers \(ATBL and ATBU\) on page 123](#). The alternate time base counter (ATB) is formed by concatenating the upper and lower alternate time base registers (ATBU and ATBL). ATBL (SPR 526) provides read-only access to the 64-bit alternate time base counter, which is incremented at an implementation-defined frequency. ATB registers are accessible in both user and supervisor mode.

Like the TB implementation, the ATBL register is an aliased name for ATB.

## 9 Storage-related APUs

This chapter describes the following APUs that are defined as part of the EIS storage architecture:

- [Section 9.1: Cache line locking APU](#)
- [Section 9.2: Direct cache flush APU](#)
- [Section 9.3: Cache way partitioning APU](#)

### 9.1 Cache line locking APU

The cache line locking APU defines instructions and methods for locking frequently used instructions and data into their cache lines. Cache locking allows software to mark individual cache lines (blocks) as locked, instructing the cache to keep latency-sensitive data available for fast access.

Unlike normal cache lines, locked cache lines do not participate in the normal replacement policy.

#### 9.1.1 Programming model

This section gives a general description of the instructions defined by the cache line locking APU. Full descriptions are provided in [Chapter 7: Instruction set](#).

##### 9.1.1.1 Lock setting and clearing

Lines are locked into the cache by software using a series of touch and lock set instructions. The following instructions are provided to lock data items into the data and instruction cache:

- **dcbtl**s—Data Cache Block Touch and Lock Set
- **dcbtstl**s—Data Cache Block Touch for Store and Lock Set
- **icbtl**s—Instruction Cache Block Touch and Lock Set

The **rA** and **rB** operands to these instructions form an effective address identifying the line to be locked. The **CT** field indicates which cache in the cache hierarchy should be targeted. These instructions are similar to the **dcbt**, **dcbtst**, and **icbt** instructions, but locking instructions can not execute speculatively and may cause additional exceptions. For unified caches, both the instruction lock set and the data lock set target the same cache.

Similarly, lines are unlocked from the cache by software using a series of lock-clear instructions. The following instructions are provided to lock instructions into the instruction cache:

- **dcblc**—Data Cache Block Lock Clear
- **icblc**—Instruction Cache Block Lock Clear

The **rA** and **rB** operands to these instructions form an EA identifying the line to be unlocked. The **CT** field indicates which cache in the cache hierarchy should be targeted.

Additionally, software may clear all the locks in the cache. For the primary cache, this is accomplished by setting the **CLFC** (**DCLFC**, **ICLFC**) bit in **L1CSR0** (**L1CSR1**).

Cache lines can also be implicitly unlocked in the following ways:

- A locked line is invalidated if it is targeted by a **dcbi**, **dcbf**, or **icbi** instruction.
- A snoop hit on a locked line that requires the line to be invalidated. This can occur because the data the line contains has been modified external to the processor, or another processor has explicitly invalidated the line.
- The entire cache containing the locked line is flash invalidated.

An implementation is not required to unlock lines if data is invalidated in the cache. Although the data may be invalidated (and thus not in the cache), the line can remain locked and be filled from the memory subsystem when the next access occurs. This method of not clearing locks when the associated line is invalidated, is called persistent locking. An implementation may choose to implement locks as persistent or not persistent; the preferred method is persistent.

### 9.1.1.2 Error conditions

Setting locks in the cache can fail for several reasons. An address specified with a lock set instruction that does not have the proper permission causes a data storage interrupt (DSI). Cache locking addresses are always translated as data references, therefore **icbtls** instructions that fail to translate or fail permissions cause DTLB and DSI errors respectively. Additionally, cache locking and clearing operations can fail due to restricted user mode access. See [Section 9.1.1.4.1: Cache locking \(user mode\) exceptions](#).

### 9.1.1.3 Overlocking

If no exceptions occur for the execution of an **dcbtls**, **dcbtstls**, or **icbtls** instruction an attempt is made to lock the corresponding line in the cache. If all of the available ways are already locked in the given cache set, the requested line is not locked. This is considered an overlocking situation and if the lock was targeted for the primary cache (CT = 0) then L1CSR0[DCLO] (or L1CSR1[ICLO] if **icbtls**) is set appropriately.

A processor may optionally allow victimizing a locked line in an overlocking situation. If L1CSR0[DCLOA] (L1CSR0[ICLOA] for the primary instruction cache,) is set, an overlocking condition causes the replacement of an existing locked line with the requested line. The selection of the line to replace in an overlocking situation is implementation dependent. The overlocking condition is still said to exist and is appropriately reflected in the status bits for lock overflow.

An attempt to lock a line that is present and valid in the cache does not cause an overlocking condition.

A non-lock-setting cache-line fill or line replacement request to a cache that has all ways locked for a given set does not cause a lock to be cleared.

### 9.1.1.4 Unable-to-lock conditions

If no exceptions occur and no overlocking condition exists, an attempt to set a lock can fail if any of the following is true:

- The target address is marked cache-inhibited or the storage attributes of the address uses a coherency protocol that does not support locking.
- The target cache is disabled or not present.
- The CT field specifies a value not supported by the implementation.
- Any other implementation-specific error condition.

If an unable-to-lock condition occurs, the lock set instruction is treated as a NOP. If the lock targeted the data cache (**dcbtls**, **dcbtstls**), L1CSR0[DCUL] is set to indicate the unable-to-lock condition; if the lock targeted the instruction cache (**icbtls**), L1CSR1[ICUL] is set. L1CSR0[DCUL] or L1CSR0[ICUL] is set regardless of the CT value in the lock-setting instruction.

#### 9.1.1.4.1 Cache locking (user mode) exceptions

Setting and clearing cache locks can be restricted to supervisor mode only access. If set, MSR[UCLE] allows cache locking operations to be performed in user mode. If MSR[UCLE] = 0 and MSR[PR] = 1 and execution of a cache lock or cache clear instruction occurs, a cache locking exception occurs. In this case the processor suppresses execution of the instruction causing the exception. A DSI interrupt is taken and SRR0, SRR1, MSR, and ESR 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[CE,ME,DE] are unchanged. All other bits are cleared.
- ESR[DLK] is set if the instruction was a **dcbtls**, **dcbtstls**, or a **dcblc**.
- ESR[ILK] is set if the instruction was a **icbtls** or a **icblc**.
- All other ESR bits are cleared.

Instruction execution resumes at address IVPR[0–47]||IVOR2[48–59]||0b0000.

## 9.2 Direct cache flush APU

### 9.2.1 Overview

To assist in software flush of the L1 cache, the direct cache flush APU allows the programmer to flush and/or invalidate the cache by specifying the cache set and cache way. Without such a feature, the programmer must either:

- Know the virtual addresses of the lines that need to be flushed and issue **dcbst** or **dcbf** instructions to those addresses.
- Flush the entire cache by causing all the lines to be replaced. This requires a virtual address range that is mapped as a contiguous physical address range, that the programmer knows and can manipulate the replacement policy of the cache, and the size and organization of the cache.

With the direct cache flush APU the program needs only specify the way and set of the cache to flush.

The direct cache flush APU available bit, L1CFG0[CFISWA], is set for implementations that contain the direct cache flush APU.

### 9.2.2 Programming model

To address a specific physical block of the cache, the L1 flush and invalidate control register 0 (L1FINV0) is written with the cache set (L1FINV0[CSET]) and cache way (L1FINV0[CWAY]) of the line that is to be flushed. L1FINV0 is written using a **mtspr** instruction specifying the L1FINV0 register. No tag match in the cache is required. An additional field, L1FINV0[CCMD], is used to specify the type of flush to be performed on the line addressed by L1FINV0[CWAY] and L1FINV0[CSET].

The available L1FINV0[CCMD] encodings are described in [Table 33 on page 99](#).

Only the L1 data cache (or unified cache) is manipulated by the direct cache flush APU. The L1 instruction cache or any other caches in the cache hierarchy are not explicitly targeted by this APU.

### 9.2.2.1 Register model

The direct cache flush APU defined one register, the L1 flush and invalidate control register 0, described in [Section 3.11.5: L1 flush and invalidate control register 0 \(L1FINV0\)](#). L1FINV0 contains fields to provide the way and set selection of a cache line to flush and or invalidate.

## 9.3 Cache way partitioning APU

The cache way partitioning APU allows ways in a unified L1 cache to be configured to accept either data or instruction miss line-fill replacements.

### 9.3.1 Programming model

The cache way partitioning APU is comprised of bits in L1CSR0 and L1CFG0, as follows:

- Way instruction disable field (L1CSR0[WID]) is a 4-bit field that determines which of ways 0–3 are available for replacement by instruction miss line refills.
- The additional ways instruction disable bit (L1CSR0[AWID]) determines whether ways 4 and above are available for replacement by instruction miss line refills.
- Way data disable field (L1CSR0[WDD]) is a 4-bit field that determines which of ways 0–3 are available for replacement by data miss line refills.
- The additional ways data disable bit (L1CSR0[AWDD]) determines whether ways 4 and above are available for replacement by instruction miss line refills.
- See [Section 3.11.1: L1 cache control and status register 0 \(L1CSR0\)](#).
- Way access mode bit, L1CSR0[WAM], Determines whether all ways are available for access or only ways partitioned for the specific type of access are used for a fetch or read operation. See [Section 3.11.1: L1 cache control and status register 0 \(L1CSR0\)](#).
- Cache way partitioning APU available bit, L1CFG0[CWPA], indicates whether the cache way partitioning APU is available. See [Section 3.11.3: L1 cache configuration register 0 \(L1CFG0\)](#).

These fields are described in detail in [Section 3.11.3: L1 cache configuration register 0 \(L1CFG0\)](#), and in [Section 3.11.1: L1 cache control and status register 0 \(L1CSR0\)](#).

### 9.3.2 Interaction with the cache locking APU

Note that the cache way partitioning APU can affect the cache line locking APU's ability to control replacement of lines. If any cache line locking instruction (**icbtlis**, **dcbtlis**, **dcbtstls**) is allowed to execute and finds a matching line in the cache, the line's lock bit is set regardless of the L1CSR0[WID,AWID,WDD,AWDD] settings. In this case, no replacement has been made.

However, for cache misses that occur while executing a cache line lock set instruction, the only candidate lines available for locking are those that correspond to ways of the cache that have not been disabled for the particular type of line locking instruction (controlled by WDD and AWDD for **dcbtlis** and **dcbtstls**, controlled by WID and AWID for **icbtlis**). Thus, an

overlocking condition may result even though fewer than eight lines with the same index are locked.



## 10 VLE introduction

This body of this document describes the VLE (variable length encoding) extension to the Book E architecture. The VLE extension offers more efficient binary representations of applications for the embedded processor spaces where code density plays a major role in affecting overall system cost, and to a somewhat lesser extent, performance. The intent of the VLE extension is not to define an entirely different ISA nor to supplant the PowerPC ISA; instead the VLE extension can be viewed as a supplement that is can be applied to an application or to part of an application to improve code density.

[Chapter 12: VLE compatibility with the EIS on page 701](#), describes additional VLE extensions to the EIS.

The major objectives of the VLE extension are as follows:

- Coexistence and consistency with the Book E ISA and general architecture
- Maintain a common programming model and instruction operation model in the VLE extension
- Reduce overall code size by ~30% over existing PowerPC text segments
- Limit the increase in execution path length to under 10% for most important applications
- Limit the increase in hardware complexity for implementations containing the VLE extension

### 10.1 Compatibility with PowerPC Book E

VLE provides an extension to Book E. There are additional operations defined using an alternate instruction encoding to enable reduced code footprint. This alternate encoding set is selected on an instruction page basis. A single page attribute bit selects between standard Book E instruction encodings and VLE instructions for that page of memory. This attribute is an extension to the Book E page attributes. Pages can be freely intermixed, allowing for a mixture of both types of encodings.

Instruction encodings in pages marked as using the VLE extension are either 16 or 32 bits long, and are aligned on 16-bit boundaries. Because of this, all instruction pages marked as VLE are required to use big-endian byte ordering.

The programmer's model uses the same register set with both instruction encodings, although certain registers are not accessible by VLE instructions using the 16-bit formats and not all condition register (CR) fields are used by condition setting or conditional branch instructions executing from a VLE instruction page. In addition, immediate fields and displacements differ in size and use, due to the more restrictive encodings imposed by VLE instructions.

The VLE extension defines additional fields in registers defined by Book E and the EIS. These are described in [Section 12.2: VLE extension processor and storage control extensions](#).

Other than the requirement of big-endian byte ordering for instruction pages and the additional page attribute to identify whether the instruction page corresponds to a VLE section of code, VLE complies with the memory model defined in Book E and the Book E Implementation Specifications (EIS). Likewise, the VLE extension complies with the Book E

and EIS definitions of the exception and interrupt model, the timer facilities, the debug facilities and the special-purpose registers (SPRs).

## 10.2 Instruction mnemonics and operands

The description of each instruction includes the mnemonic and a formatted list of operands. VLE instruction semantics are either identical or similar to Book E instruction semantics. Where the semantics, side-effects, and binary encodings are identical, Book E mnemonics and formats are used. Where the semantics are similar but the binary encodings differ, the Book E mnemonic is typically preceded with an **e\_**. To distinguish similar instructions available in both 16- and 32-bit forms under VLE and standard Book E instructions, VLE instructions encoded with 16 bits have an **se\_** prefix. Those VLE instructions encoded with 32 bits that have different binary encodings or semantics than the equivalent Book E instruction have an **e\_** prefix. The following are examples:

```
stw rS,D(rA)           // standard Book E instruction
e_stw rS,D(rA)         // 32-bit VLE instruction
se_stw rZ,SD4(rX)     // 16-bit VLE instruction
```

## 11 VLE storage addressing

A program references memory using the effective address (EA) computed by the processor when it executes a branch, storage access, storage control, or TLB management instruction, or when it fetches the next sequential instruction.

### 11.1 Data memory addressing modes

[Table 218](#) lists data memory addressing modes supported by the VLE extension.

**Table 218. Data storage addressing modes**

Mode	Name	Description
Base+16-bit displacement (32-bit instruction format)	D-mode	The 16-bit D field is sign-extended and added to the contents of the GPR designated by rA or to zero if rA = 0 to produce the EA.
Base+8-bit displacement (32-bit instruction format)	D8-mode	The 8-bit D8 field is sign-extended and added to the contents of the GPR designated by rA or to zero if rA = 0 to produce the EA.
Base+scaled 4-bit displacement (16-bit instruction format)	SD4-mode	The 4-bit SD4 field zero-extended, scaled (shifted left) according to the size of the operand, and added to the contents of the GPR designated by rX to produce the EA. (Note that rX = 0 is not a special case).
Base+Index (32-bit instruction format)	X-mode	The GPR contents designated by rB are added to the GPR contents designated by rA or to zero if rA = 0 to produce the EA.

### 11.2 Instruction memory addressing modes

[Table 219](#) lists instruction memory addressing modes supported by the VLE extension.

**Table 219. Instruction storage addressing modes**

Mode	Description
I-form branch instructions (32-bit instruction format)	The 24-bit BD24 field is concatenated on the right with 0b0, sign-extended, and then added to the address of the branch instruction.
Taken B15-form branch instructions (32-bit instruction format)	The 15-bit BD15 field is concatenated on the right with 0b0, sign-extended, and then added to the address of the branch instruction to form the EA of the next instruction.
All branch instructions (16-bit instruction format)	The 8-bit BD8 field is concatenated on the right with 0b0, sign-extended, and then added to the address of the branch instruction to form the EA of the next instruction.
Sequential instruction fetching (or non-taken branch instructions)	The value 4 [2] is added to the address of the current 32-bit [16-bit] instruction to form the EA of the next instruction. If the address of the current instruction is 0xFFFF_FFFC [0xFFFF_FFFE], the address of the next sequential instruction is undefined.

Table 219. Instruction storage addressing modes (continued)

Mode	Description
Any branch instruction with LK = 1 (32-bit instruction format)	The value 4 is added to the address of the current branch instruction and the result is placed into the LR. If the address of the current instruction is 0xFFFF_FFFC, the result placed into the LR is undefined.
Branch <b>se_bl</b> , <b>se_birl</b> , <b>se_bctrl</b> instructions (16-bit instruction format)	The value 2 is added to the address of the current branch instruction and the result is placed into the LR. If the address of the current instruction is 0xFFFF_FFFE, the result placed into the LR is undefined.

## 12 VLE compatibility with the EIS

The body of this document addresses the relationship between VLE and Book E. It does not explicitly address EIS-defined features, such as the APUs or the use of MAS registers. However, the information in the previous chapters provides a model for how the VLE extension is integrated with features defined by the layer of architecture defined by the EIS.

### 12.1 Overview

The VLE extension uses the same semantics as the Book E architecture. Due to the limited instruction encoding formats, VLE instructions typically support reduced immediate fields and displacements, and not all Book E operations are encoded in the VLE extension. The basic philosophy is to capture all useful operations, with most frequent operations given priority. Immediate fields and displacements are provided to cover the majority of ranges encountered in embedded control code. Instructions are encoded in either a 16- or 32-bit format, and these may be freely intermixed.

Book E floating-point registers (FPRs) are not accessible by VLE instructions. VLE instructions use Book E GPR and SPR registers with the following limitations:

- VLE instructions using the 16-bit formats are limited to addressing GPR0–GPR7, and GPR24–GPR31 in most instructions. Move instructions are provided to transfer register contents between these registers and GPR8–GPR23.
- VLE instructions using the 16-bit formats are limited to addressing CR0
- VLE instructions using the 32-bit formats are limited to addressing CR0–CR3

VLE instruction encodings are generally different than Book E instructions, except that most Book E instructions falling within Book E major opcode 31 are encoded identically in 32-bit VLE instructions and have identical semantics unless they affect or access a resource not supported by the VLE extension. Also, major opcode 4 is available to support additional APUs using identical encodings for both Book E and the VLE extension. This allows an implementation of the VLE extension to include additional APUs, such as the cache-line locking, single-precision floating-point, and SPE APUs, and to use the exact encodings.

Because future compatibility is desired, and to avoid confusion with Book E, register bit numbering remains the same as in Book E.

### 12.2 VLE extension processor and storage control extensions

This section describes additional functionality and extensions to the EIS to support the VLE extension.

#### 12.2.1 EIS instruction extensions

This section describes extensions to EIS instructions to support VLE operations. Because instructions may reside on a half-word boundary, bit 62 is not masked by instructions that cause fetching from a register, such as the LR, CTR, or a save/restore register 0, that holds an instruction address:

- Return from interrupt instructions, such as **rfdi** (defined as part of the debug APU) and **rfmci** (defined as part of the machine check APU) no longer mask bit 62 of the respective save/restore register 0. The destination address is  $xSRR0[32-62] \parallel 1'b0$ .

### 12.2.2 Book E instruction extensions

This section describes the various extensions to Book E instructions to support the VLE extension:

- **rfdi**, **rfdi**, and **rfdi** no longer mask bit 62 of CSRR0, DSRR0, or SRR0. The destination address is  $xSRR0[32-62] \parallel 1'b0$ .
- **bclr**, **bclrl**, **bcctr**, and **bcctrl** no longer mask bit 62 of the LR or CTR. The destination address is  $[LR,CTR][32-62] \parallel 1'b0$ .

### 12.2.3 EIS MMU extensions

The VLE assumes that the MMU implementation complies with the more general MMU definition provided by Book E and the more specific definition provided by the EIS. This section describes the differences and extensions to the MMU necessary to support the VLE extension.

#### 12.2.3.1 TLB entries

Each TLB entry is augmented with an additional page attribute bit, the VLE bit. If set, VLE indicates the corresponding page of memory is a VLE page.

#### 12.2.3.2 TLB load on reset

During reset, all TLB entries except entry 0 are invalidated. TLB entry 0 is loaded with the additional value shown in [Table 220](#).

**Table 220. TLB Entry 0 reset value**

Field	Reset value	Comments
VLE	<i>p_rst_vlemode</i> value	Book E mode, not VLE if no <i>p_rst_vlemode</i> signal is available

Note that implementations may provide a *p\_rst\_vlemode* input to supply the value of the VLE field on reset. If not available, the default value should be 0, indicating a Book E page

#### 12.2.3.3 VLE attribute bit


If set, the VLE attribute bit indicates the corresponding page of memory is a VLE page. The VLE attribute is used only for instruction access and is ignored for data accesses. The VLE bit may be set only for big-endian pages, otherwise a byte-ordering exception occurs on instruction fetches.

#### 12.2.3.4 MMU assist registers (MAS<sub>n</sub>)

To support the VLE extension, additional bits are defined in MAS2 and MAS4. These are described in the following sections.

##### MAS2

The MAS2 register is shown below. The VLE page attribute has been added as MAS2[58]. If the VLE extension is not present, this bit is always read as zero and writes are ignored.




MAS2[VLE] is defined in [Table 221](#).

**Table 221. MAS2 field descriptions**

Bits	Name	Comments, or function when set
58	VLE	VLE 0 This page is a standard Book E page 1 This page is a VLE page

MMU assist register 4 (MAS4)

When the VLE extension is implemented, MAS4[58] is defined as the VLED field, which contains the default MAS2[VLE] value. If the VLE extension is not present, this bit is always read as zero and writes are ignored. MAS4 is shown below.



MAS4[VLED] is described in [Table 222](#).

**Table 222. MAS4 field descriptions**

Bits	Name	Comments, or function when set
58	VLED	Default VLE value. Defined by the EIS. 0 This page is a standard Book E page 1 This page is a VLE page

### 12.2.4 EIS debug APU extensions

The **se\_rfdi** instruction is provided to support the EIS debug interrupt APU.

**rfdi** **rfdi**

#### Return from debug interrupt

**se\_rfdi**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0

MSR ← DSRR1

NIA ← DSRR0<sub>32:62</sub> || 0b0

The **se\_rfdi** instruction is used to return from a debug class interrupt, or as a means of establishing a new context and synchronizing on that new context simultaneously.

The contents of DSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address DSRR0[32–62]||0b0. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism (see Book E) is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in DSRR0 at the time of the execution of **se\_rfdi**).

Execution of this instruction is privileged and restricted to supervisor mode only.

Execution of this instruction is context synchronizing.

When the debug APU is disabled, this instruction is treated as an illegal instruction.

Special Registers Altered: MSR



## 13 VLE instruction classes

This chapter lists instructions defined or supported by the VLE extension. Unless otherwise noted, instructions that are not prefixed with **e\_** or **se\_** have identical encodings and semantics as in Book E or in the Book E implementation standards (EIS). Full descriptions of these instructions are provided in the EREF: Programmers reference manual for ST's Book E processors.

A complete list of supported instructions is provided in [Section 13.6: Instruction listings](#).

### 13.1 Processor control instructions

This section lists processor control instructions that can be executed when a processor is in VLE mode. These instructions are grouped as follows:

- [Section 13.1.1: System linkage instructions](#)
- [Section 13.1.2: Processor control register manipulation instructions](#)
- [Section 13.1.3: Instruction synchronization instruction](#)

#### 13.1.1 System linkage instructions

**se\_sc**, **se\_rfi**, **se\_rfci**, and **se\_rfdi** are system linkage instructions that enable the program to call upon the system to perform a service (that is, invoke a system call interrupt), and by which the system can return from performing a service or from processing an interrupt.

[Table 223](#) lists system linkage instructions.

**Table 223. System linkage instruction set index**

Mnemonic	Instruction	Reference
<b>se_sc</b>	System Call	<a href="#">on page 769</a>
<b>se_rfci</b>	Return From Critical Interrupt	<a href="#">on page 767</a>
<b>se_rfdi</b>	Return From Debug Interrupt	<a href="#">on page 704</a>
<b>se_rfi</b>	Return From Interrupt	<a href="#">on page 767</a>

#### 13.1.2 Processor control register manipulation instructions

In addition to the Book E processor control register manipulation instructions, the VLE extension provides 16-bit forms of instructions to move to/from the LR and CTR. [Table 224](#) lists the processor control register manipulation instructions.

**Table 224. System register manipulation instruction set index**

Mnemonic	Instruction	Reference
<b>se_mfctr</b> rX	Move From Count Register	<a href="#">on page 762</a>
<b>mfocr</b> rD,DCRN	Move From Device Control Register	Book E
<b>se_mflr</b> rX	Move From Link Register	<a href="#">on page 762</a>
<b>mfmsr</b> rD	Move From Machine State Register	Book E

**Table 224. System register manipulation instruction set index (continued)**

Mnemonic	Instruction	Reference
<b>mfspir</b> rD,SPRN	Move From Special Purpose Register	Book E
<b>se_mtctr</b> rX	Move To Count Register	<a href="#">on page 763</a>
<b>mtdcr</b> DCRN,rS	Move To Device Control Register	Book E
<b>se_mtlr</b> rX	Move To Link Register	<a href="#">on page 763</a>
<b>mtmsr</b> rS	Move To Machine State Register	Book E
<b>mtspr</b> SPRN,rS	Move To Special Purpose Register	Book E
<b>wrttee</b> rA	Write MSR External Enable	Book E
<b>wrtteei</b> E	Write MSR External Enable Immediate	Book E

### 13.1.3 Instruction synchronization instruction

[Table 225](#) lists the VLE-defined **se\_isync** instruction.

**Table 225. Instruction Synchronization Instruction Set Index**

Mnemonic	Instruction	Reference
<b>se_isync</b>	Instruction Synchronize	<a href="#">on page 757</a>

## 13.2 Branch operation instructions

This section lists branch instructions that can be executed when a processor is in VLE mode and the registers that support them.

### 13.2.1 Registers for branch operations

The registers that support branch operations are grouped as follows:

- [Section 3.5.1: Condition register \(CR\)](#)
- [Section 3.5.2: Link register \(LR\)](#)
- [Section 3.5.3: Count register \(CTR\)](#)

#### 13.2.1.1 Condition register (CR)

The condition register (CR) is a 32-bit register. CR bits are numbered 32 (most-significant bit) to 63 (least-significant bit). The CR reflects the result of certain operations, and provides a mechanism for testing (and branching). The VLE extension implements the entire CR, but some comparison operations and all branch instructions are limited to using CR0–CR3. The full Book E condition register field and logical operations are provided however.

Access: User read/write

31 30 29 28 | 27 26 25 24 | 23 22 21 20 | 19 18 17 16 | 15 14 13 12 | 11 10 9 8 | 7 6 5 4 | 3 2 1 0

R	CR0	CR1	CR2	CR3	CR4	CR5	CR6	CR7
W								
Reset	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0	0 0 0 0

CR bits are grouped into eight 4-bit fields, CR0–CR7, which are set in one of the following ways.

- Specified fields of the condition register can be set by a move to the CR from a GPR (**mtcrf**).
- A specified CR field can be set by a move to the CR from another CR field (**e\_mcrf**).
- CR field 0 can be set as the implicit result of an integer instruction.
- A specified condition register field can be set as the result of an integer compare instruction.
- CR field 0 can be set as the result of an integer bit test instruction.

Instructions are provided to perform logical operations on individual CR bits and to test individual condition register bits (see Book E).

### 13.2.1.2 Condition register settings for integer instructions

For all integer word instructions in which the Rc bit is defined and set, and for **addic.**, the first three bits of CR field 0 (CR[32–34]) are set by signed comparison of bits 32–63 of the result to zero, and the fourth bit of CR field 0 (CR[35]) is copied from the final state of XER[SO].

```

if (target_register)32:63 < 0 then c ← 0b100
else if (target_register)32:63 > 0 then c ← 0b010
else c ← 0b001
CR0 ← c || XERSO
    
```

If any portion of the result is undefined, the value placed into the first three bits of CR field 0 is undefined.

The bits of CR field 0 are interpreted as shown in [Table 226](#).

**Table 226. CR0 encodings**

CR Bit	Description
32	Negative (LT). Bit 32 of the result is equal to 1.
33	Positive (GT). Bit 32 of the result is equal to 0 and at least one of bits 33–63 of the result is non-zero.
34	Zero (EQ). Bits 32–63 of the result are equal to 0.
35	Summary overflow (SO). This is a copy of the final state XER[SO] at the completion of the instruction.

### 13.2.1.3 Condition register setting for compare instructions

For compare instructions, a CR field specified by the **crD** operand in for the **e\_cmph**, **e\_cmphl**, **e\_cmpi**, and **e\_cmpli** instructions, or CR0 for the **e\_cmp16i**, **e\_cmph16i**, **e\_cmphl16i**, **e\_cmpl16i**, **se\_cmp**, **se\_cmph**, **se\_cmphl**, **se\_cmpi**, and **se\_cmpli** instructions, is set to reflect the result of the comparison. The CR field bits are interpreted as shown in [Table 227](#). A complete description of how the bits are set is given in the instruction

descriptions and [Section 13.4.5: Integer compare and bit test instructions](#).

**Table 227. Condition register setting for compare instructions**

CR Bit	Description
4×CRD + 32	Less than (LT) For signed-integer compare, GPR(rA or rX) < SCI8 or SI or GPR(rB or rY). For unsigned-integer compare, GPR(rA or rX) < <sub>u</sub> SCI8 or UI or UI5 or GPR(rB or rY).
4×CRD + 33	Greater than (GT) For signed-integer compare, GPR(rA or rX) > SCI8 or SI or UI5 or GPR(rB or rY). For unsigned-integer compare, GPR(rA or rX) > <sub>u</sub> SCI8 or UI or UI5 or GPR(rB or rY).
4×CRD + 34	Equal (EQ) For integer compare, GPR(rA or rX) = SCI8 or UI5 or SI or UI or GPR(rB or rY).
4×CRD + 35	Summary overflow (SO) For integer compare, this is a copy of the final state of XER[SO] at the completion of the instruction.

**13.2.1.4 Condition register setting for the bit test instruction**

The Bit Test Immediate instruction, **se\_btsti**, also sets CR field 0. See the instruction description and also [Section 13.4.5: Integer compare and bit test instructions](#).

**13.2.1.5 Link register (LR)**

VLE instructions use the LR as defined in Book E, although the VLE extension defines a subset of all variants of Book E conditional branches involving the LR, as shown in [Table 228](#).

**Table 228. Branch to link register instruction comparison**

Book E		VLE subset	
Instruction	Syntax	Instruction	Syntax
Branch Conditional to Link Register Branch Conditional to Link Register & Link	<b>bclr</b> BO,BI <b>bclrl</b> BO,BI	Branch (Absolute) to Link Register Branch (Absolute) to Link Register & Link	<b>se_blr</b> <b>se_blrI</b>
Branch Conditional & Link	<b>e_bcl</b> BO,BI,BD	Branch Conditional & Link	<b>e_bcl</b> BO32,BI32,BD15
		Branch (Absolute) & Link	<b>e_bl</b> BD24 <b>se_bl</b> BD8

**13.2.1.6 Count register**

VLE instructions use the count register (CTR) as defined in Book E, although the VLE extension defines a subset of the variants of Book E conditional branches involving the CTR, as shown in [Table 229](#).

**Table 229. Branch to count register instruction comparison**

Book E		VLE	
Instruction	Syntax	Instruction	Syntax
Branch Conditional to Count Register Branch Conditional to Count Register & Link	<b>bcctr</b> BO,BI <b>bcctrl</b> BO,BI	Branch (Absolute) to Count Register Branch (Absolute) to Count Register & Link	<b>se_bctr</b> <b>se_bctrl</b>

### 13.2.2 Branch instructions

The sequence of instruction execution can be changed by the branch instructions. Because VLE instructions must be aligned on half-word boundaries, the low-order bit of the generated branch target address is forced to 0 by the processor in performing the branch.

The branch instructions compute the EA of the target in one of the following ways, as described in [Section 11.2: Instruction memory addressing modes](#).

1. Adding a displacement to the address of the branch instruction.
2. Using the address contained in the LR (Branch to Link Register [and Link]).
3. Using the address contained in the CTR (Branch to Count Register [and Link]).

Branching can be conditional or unconditional, and the return address can optionally be provided. If the return address is to be provided (LK = 1), the EA of the instruction following the branch instruction is placed into the LR after the branch target address has been computed: this is done whether or not the branch is taken.

In branch conditional instructions, the BI32 or BI16 instruction field specifies the CR bit to be tested. For 32-bit instructions using BI32, CR[32–47] (corresponding to bits in CR0–CR3) may be specified. For 16-bit instructions using BI16, only CR[32–35] (bits within CR0) may be specified.

In branch conditional instructions, the BO32 or BO16 field specifies the conditions under which the branch is taken and how the branch is affected by or affects the CR and CTR. Note that VLE instructions also have different encodings for the BO32 and BO16 fields than in Book E’s BO field.

If the BO32 field specifies that the CTR is to be decremented, CTR[32–63] are decremented. If BO[16,32] specifies a condition that must be TRUE or FALSE, that condition is obtained from the contents of CR[BI+32]. (Note that CR bits are numbered 32–63. BI refers to the BI field in the branch instruction encoding. For example, specifying BI = 2 refers to CR[34].)

Encodings for the BO32 field for the VLE extension are shown in [Table 230](#).

**Table 230. VLE extension BO32 encodings**

BO32	Description
00	Branch if the condition is FALSE.
01	Branch if the condition is TRUE.
10	Decrement CTR[32–63], then branch if the decremented CTR[32–63]≠0.
11	Decrement CTR[32–63], then branch if the decremented CTR[32–63] = 0.

The encoding for the BO16 field for the VLE extension is shown in [Table 231](#).

**Table 231. VLE extension BO16 encodings**

BO16	Description
0	Branch if the condition is FALSE.
1	Branch if the condition is TRUE.

The various branch instructions supported by the VLE extension are shown in [Table 232](#).

**Table 232. Branch instruction set index**

Mnemonic	Instruction	Reference
<b>e_b</b> BD24 <b>e_bl</b> BD24	Branch Branch & Link	<a href="#">on page 744</a>
<b>se_b</b> BD8 <b>se_bl</b> BD8	Branch Branch & Link	<a href="#">on page 744</a>
<b>e_bc</b> BO32,BI32,BD15 <b>se_bc</b> BO16,BI16,BD8 <b>e_bcl</b> BO32,BI32,BD15	Branch Conditional Branch Conditional Branch Conditional & Link	<a href="#">on page 745</a>
<b>se_bctr</b> <b>se_bctrl</b>	Branch to Count Register Branch to Count Register & Link	<a href="#">on page 746</a>
<b>se_blr</b> <b>se_blrl</b>	Branch to Link Register Branch to Link Register & Link	<a href="#">on page 747</a>

### 13.3 Condition register instructions

Condition register instructions are provided to transfer values to/from various portions of the CR. The VLE extension does not introduce any additional functionality beyond that defined in Book E for CR operations, but does remap the CR-logical and **mcrf** instruction functionality into major opcode 31. These instructions operate identically to the Book E instructions, but are encoded differently. [Table 233](#) lists condition register instructions supported in VLE mode.

**Table 233. Condition register instruction set index**

Mnemonic	Instruction	Reference
<b>e_crand</b> crbD,crbA,crbB	Condition Register AND	<a href="#">on page 753</a>
<b>e_crandc</b> crbD,crbA,crbB	Condition Register AND with Complement	<a href="#">on page 753</a>
<b>e_creqv</b> crbD,crbA,crbB	Condition Register Equivalent	<a href="#">on page 753</a>
<b>e_crnand</b> crbD,crbA,crbB	Condition Register NAND	<a href="#">on page 754</a>
<b>e_crnor</b> crbD,crbA,crbB	Condition Register NOR	<a href="#">on page 754</a>
<b>e_cror</b> crbD,crbA,crbB	Condition Register OR	<a href="#">on page 755</a>
<b>e_crorc</b> crbD,crbA,crbB	Condition Register OR with Complement	<a href="#">on page 755</a>
<b>e_crxor</b> crbD,crbA,crbB	Condition Register XOR	<a href="#">on page 755</a>

Table 233. Condition register instruction set index

Mnemonic	Instruction	Reference
<b>e_mcrf</b> crD,crS	Move Condition Register Field	<a href="#">on page 761</a>
<b>mcrxr</b> crD	Move to Condition Register from Integer Exception Register	Book E
<b>mfcrr</b> rD	Move From condition register	Book E
<b>mtcrf</b> FXM,rS	Move to Condition Register Fields	Book E

## 13.4 Integer instructions

This section lists the integer instructions supported by the VLE extension.

### 13.4.1 Integer load instructions

The integer load instructions compute the EA of the memory to be accessed as described in [Section 11.1: Data memory addressing modes](#).

The byte, half word, or word in memory addressed by EA is loaded into GPR(rD) or GPR(rZ).

The VLE extension supports both big- and little-endian byte ordering for data accesses.

Some integer load instructions have an update form in which GPR(rA) is updated with the EA. For these forms, if  $rA \neq 0$  and  $rA \neq rD$ , the EA is placed into GPR(rA) and the memory element (byte, half word, word, or double word) addressed by EA is loaded into GPR(rD). If  $rA = 0$  or  $rA = rD$ , the instruction form is invalid. This is the same behavior as specified for load with update instructions in Book E.

Basic integer load instructions are listed in [Table 234](#).

Table 234. Basic integer load instruction set index

Mnemonic	Instruction	Reference
<b>e_lbz</b> rD,D(rA) <b>e_lbzu</b> rD,D8(rA) <b>se_lbz</b> rZ,SD4(rX)	Load Byte and Zero Load Byte and Zero with Update Load Byte and Zero (16-bit form)	<a href="#">on page 757</a>
<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	Book E
<b>e_lha</b> rD,D(rA) <b>e_lhau</b> rD,D8(rA)	Load Halfword Algebraic Load Halfword Algebraic with Update	<a href="#">on page 758</a>
<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	Book E
<b>e_lhz</b> rD,D(rA) <b>e_lhzu</b> rD,D8(rA) <b>se_lhz</b> rZ,SD4(rX)	Load Halfword and Zero Load Halfword and Zero with Update Load Halfword and Zero (16-bit form)	<a href="#">on page 759</a>
<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	Book E

**Table 234. Basic integer load instruction set index (continued)**

Mnemonic	Instruction	Reference
<b>e_lwz</b> rD,D(rA)	Load Word and Zero	<a href="#">on page 761</a>
<b>e_lwzu</b> rD,D8(rA)	Load Word and Zero with Update	
<b>se_lwz</b> rZ,SD4(rX)	Load Word and Zero (16-bit form)	
<b>lwzx</b> rD,rA,rB	Load Word and Zero Indexed	Book E
<b>lwzux</b> rD,rA,rB	Load Word and Zero with Update Indexed	

Integer load byte-reversed instructions are listed in [Table 235](#).

**Table 235. Integer Load Byte-Reverse Instruction Set Index**

Mnemonic	Instruction	Reference
<b>lbrx</b> rD,rA,rB	Load Halfword Byte-Reverse Indexed	Book E
<b>lwbx</b> rD,rA,rB	Load Word Byte-Reverse Indexed	Book E

The VLE-defined integer load multiple instruction is listed in [Table 236](#).

**Table 236. Integer load multiple instruction set index**

Mnemonic	Instruction	Reference
<b>e_lmw</b> rD,D8(rA)	Load Multiple Word	<a href="#">on page 760</a>

The VLE-defined integer load and reserve instruction is listed in [Table 237](#).

**Table 237. Integer load and reserve instruction set index**

Mnemonic	Instruction	Reference
<b>lwarx</b> rD,rA,rB	Load Word And Reserve Indexed	Book E

### 13.4.2 Integer store instructions

The integer store instructions compute the EA of the memory to be accessed as described in [Section 11.1: Data memory addressing modes](#).

The contents of GPR(rS) or GPR(rZ) are stored into the byte, half word, or word in memory addressed by EA.

The VLE extension supports both big- and little-endian byte ordering for data accesses.

Some integer store instructions have an update form, in which GPR(rA) is updated with the EA. For these forms, the following rules (from Book E) apply.

- If  $rA \neq 0$ , the EA is placed into GPR(rA).
- If  $rS = rA$ , the contents of GPR(rS) are copied to the target memory element and then EA is placed into GPR(rA).

The basic integer store instructions are listed in [Table 238](#).



**Table 238. Basic integer store instruction set index**

Mnemonic	Instruction	Reference
<b>e_stb</b> rS,D(rA) <b>e_stbu</b> rS,D8(rA) <b>se_stb</b> rZ,SD4(rX)	Store Byte Store Byte with Update Store Byte (16-bit form)	<a href="#">on page 773</a>
<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Store Byte Indexed Store Byte with Update Indexed	Book E
<b>e_sth</b> rS,D(rA) <b>e_sthu</b> rS,D8(rA) <b>se_sth</b> rZ,SD4(rX)	Store Halfword Store Halfword with Update Store Halfword (16-bit form)	<a href="#">on page 773</a>
<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Store Halfword Indexed Store Halfword with Update Indexed	Book E
<b>e_stw</b> rS,D(rA) <b>e_stwu</b> rS,D8(rA) <b>se_stw</b> rZ,SD4(rX)	Store Word Store Word with Update Store Word (16-bit form)	<a href="#">on page 775</a>
<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Store Word Indexed Store Word with Update Indexed	Book E

The integer store byte-reverse instructions are listed in [Table 239](#).

**Table 239. Integer store byte-reverse instruction set index**

Mnemonic	Instruction	Reference
<b>sthbrx</b> rS,rA,rB	Store Halfword Byte-Reverse Indexed	Book E
<b>stwbrx</b> rS,rA,rB	Store Word Byte-Reverse Indexed	Book E

The integer store multiple instruction is listed in [Table 240](#).

**Table 240. Integer store multiple instruction set index**

Mnemonic	Instruction	Reference
<b>e_stmw</b> rS,D8(rA)	Store Multiple Word	<a href="#">on page 774</a>

The integer store conditional instruction is listed in [Table 241](#).

**Table 241. Integer store conditional instruction set index**

Mnemonic	Instruction	Reference
<b>stwcx.</b> rS,rA,rB	Store Word Conditional Indexed	Book E

### 13.4.3 Integer arithmetic instructions

The integer arithmetic instructions use the contents of the GPRs as source operands, and place results into GPRs, into status bits in the XER and into CR0.

The integer arithmetic instructions treat source operands as signed, two's complement integers unless the instruction is explicitly identified as performing an unsigned operation.

The **e\_add2i**. instruction and the OIM5-form instruction, **se\_subi**., set the first three bits of CR0 to characterize bits 32–63 of the result. These bits are set by signed comparison of bits 32–63 of the result to zero.

**e\_addic**[.] and **e\_subfic**[.] always set CA to reflect the carry out of bit 32.

The integer arithmetic instructions are listed in [Table 242](#).

**Table 242. Integer arithmetic instruction set index**

Mnemonic	Instruction	Reference
<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Add	Book E
<b>se_add</b> rX,rY	Add	<a href="#">on page 741</a>
<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Add Carrying	Book E
<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Add Extended	Book E
<b>e_addi</b> rD,rA,SCI8 <b>e_addi.</b> rD,rA,SCI8 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI <b>se_addi</b> rX,OIMM	Add Immediate	<a href="#">on page 741</a>
<b>e_addic</b> rD,rA,SCI8 <b>e_addic.</b> rD,rA,SCI8	Add Immediate Carrying	<a href="#">on page 743</a>
<b>e_add2is</b> rD,SI	Add Immediate Shifted	<a href="#">on page 741</a>
<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Divide Word	Book E
<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Divide Word Unsigned	Book E
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word	Book E
<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Multiply High Word Unsigned	Book E

Table 242. Integer arithmetic instruction set index (continued)

Mnemonic	Instruction	Reference
<b>e_mulli</b> rD,rA,SCI8 <b>e_mull2i</b> rD,SI	Multiply Low Immediate	<a href="#">on page 763</a>
<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Multiply Low Word	Book E
<b>se_mullw</b> rX,rY	Multiply Low Word	<a href="#">on page 764</a>
<b>neg</b> rD,rA <b>se_neg</b> rX <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Negate	<a href="#">on page 764</a>
<b>se_sub</b> rX,rY	Subtract	<a href="#">on page 775</a>
<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Subtract From	Book E
<b>se_subf</b> rX,rY	Subtract From	<a href="#">on page 776</a>
<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Subtract From Carrying	Book E
<b>e_subfc</b> rD,rA,SCI8 <b>e_subfc.</b> rD,rA,SCI8	Subtract From Immediate Carrying	<a href="#">on page 776</a>
<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	Subtract Immediate	<a href="#">on page 776</a>

#### 13.4.4 Integer logical and move instructions

Logical instructions perform bit-parallel operations on 32-bit operands or move register or immediate values into registers. The move instructions move values into a GP from either another GPR, or an immediate value.

The X-form logical instructions with  $Rc = 1$  and the SCI8-form logical instructions with  $Rc = 1$  set the first three bits of CR field 0 as described in [Section 13.4.3: Integer arithmetic instructions](#). The logical instructions do not change XER[SO,OV,CA].

The integer logical instructions are listed in [Table 243](#).

Table 243. Integer logical instruction set index

Mnemonic	Instruction	Reference
<b>and[.]</b> rA,rS,rB <b>se_and[.]</b> rX,rY	AND	<a href="#">on page 743</a>
<b>andc[.]</b> rA,rS,rB <b>se_andc</b> rX,rY	AND with Complement	<a href="#">on page 743</a>
<b>e_andi[.]</b> rA,rS,SCI8 <b>se_andi</b> rX,UI5 <b>e_and2i.</b> rD,UI	AND Immediate	<a href="#">on page 743</a>
<b>e_and2is.</b> rD,UI	AND Immediate Shifted	<a href="#">on page 743</a>
<b>se_bclri</b> rX,UI5	Bit Clear	<a href="#">on page 746</a>
<b>se_bgeni</b> rX,UI5	Bit Generate	<a href="#">on page 747</a>
<b>se_bmski</b> rX,UI5	Bit Mask Generate	<a href="#">on page 748</a>
<b>se_bseti</b> rX,UI5	Bit Set	<a href="#">on page 748</a>
<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Count Leading Zeros Word	Book E
<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Equivalent	Book E
<b>extsb</b> rA,rS <b>extsb.</b> rA,rS <b>se_extsb</b> rX	Extend Sign Byte	<a href="#">on page 755</a>
<b>extsh</b> rA,rS <b>extsh.</b> rA,rS <b>se_extsh</b> rX	Extend Sign Halfword	<a href="#">on page 755</a>
<b>se_extzb</b> rX	Extend with Zeros Byte	<a href="#">on page 756</a>
<b>se_extzh</b> rX	Extend with Zeros Halfword	<a href="#">on page 756</a>
<b>e_li</b> rD,LI20 <b>se_li</b> rX,UI7	Load Immediate	<a href="#">on page 760</a>
<b>e_lis</b> rD,UI	Load Immediate Shifted	<a href="#">on page 760</a>
<b>se_mfar</b> rX,arY	Move from Alternate Register	<a href="#">on page 762</a>
<b>se_mr</b> rX,rY	Move Register	<a href="#">on page 762</a>
<b>se_mtar</b> arX,rY	Move to Alternate Register	<a href="#">on page 763</a>
<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	NAND	Book E
<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	NOR	Book E
<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB <b>se_or</b> rX,rY	OR	<a href="#">on page 765</a>
<b>se_not</b> rX	NOT	<a href="#">on page 764</a>

**Table 243. Integer logical instruction set index (continued)**

Mnemonic	Instruction	Reference
<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	OR with Complement	Book E
<b>e_ori</b> [.] rA,rS,SCI8 <b>e_or2i</b> rD,UI	OR Immediate	<a href="#">on page 765</a>
<b>e_or2is</b> rD,UI	OR Immediate Shifted	<a href="#">on page 765</a>
<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	XOR	Book E
<b>e_xori</b> [.] rA,rS,SCI8	XOR Immediate	<a href="#">on page 765</a>

### 13.4.5 Integer compare and bit test instructions

The integer compare instructions compare the contents of GPR(rA) with one of the following:

- The value of the SCI8 field
- The zero-extended value of the UI field
- The zero-extended value of the UI5 field
- The sign-extended value of the SI field
- The contents of GPR(rB) or GPR(rY).

The following comparisons are signed: **e\_cmph**, **e\_cmpi**, **e\_cmp16i**, **e\_cmph16i**, **se\_cmp**, **se\_cmph**, and **se\_cmpi**.

The following comparisons are unsigned: **e\_cmphi**, **e\_cmpli**, **e\_cmphi16i**, **e\_cmpl16i**, **se\_cmphi**, **se\_cmpli**, and **se\_cmphi**.

When operands are treated as 32-bit signed quantities, GPRn[32] is the sign bit. When operands are treated as 16-bit signed quantities, GPRn[48] is the sign bit.

For 32-bit implementations, the L field must be zero.

Compare instructions set one of the left-most three bits of the designated CR field and clears the other two. XER[SO] is copied to bit 3 of the designated CR field.

The CR field is set as shown in [Table 244](#).

**Table 244. CR settings for compare instructions**

Bit	Name	Description
0	LT	(rA or rX) < SCI8, SI, UI5, or GPR(rB or rY) (signed comparison) (rA or rX) < <sub>u</sub> SCI8, UI, UI5 or GPR(rB or rY) (unsigned comparison)
1	GT	(rA or rX) > SCI8, SI, UI5, or GPR(rB or rY) (signed comparison) (rA or rX) > <sub>u</sub> SCI8, UI, UI5 or GPR(rB or rY) (unsigned comparison)
2	EQ	(rA or rX) = SCI8, SI, UI, UI5, or GPR(rB or rY)
3	SO	Summary overflow from the XER

The integer bit test instruction tests the bit specified by the UI5 instruction field and sets the CR0 field as shown in [Table 245](#).

**Table 245. CR settings for integer bit test instructions**

Bit	Name	Description
0	LT	Always cleared
1	GT	$RX_{ui5} == 1$
2	EQ	$RX_{ui5} == 0$
3	SO	Summary overflow from the XER

[Table 246](#) is an index for integer compare and bit test operations.

**Table 246. Integer compare and bit test instruction set index**

Mnemonic	Instruction	Reference
<b>se_btsti</b> rX,UI5	Bit Test Immediate	<a href="#">on page 748</a>
<b>cmp</b> crD,L,rA,rB <b>se_cmp</b> rX,rY	Compare	<a href="#">on page 749</a>
<b>e_cmph</b> crD,rA,rB <b>se_cmph</b> rX,rY	Compare Halfword	<a href="#">on page 750</a>
<b>e_cmph16i</b> rA,SI16	Compare Halfword Immediate	<a href="#">on page 750</a>
<b>e_cmphi</b> crD,rA,rB <b>se_cmphi</b> rX,rY	Compare Halfword Logical	<a href="#">on page 751</a>
<b>e_cmphi16i</b> rA,UI16	Compare Halfword Logical Immediate	<a href="#">on page 751</a>
<b>e_cmpi</b> crD,rA,SCI8 <b>e_cmp16i</b> rA,SI16 <b>se_cmpi</b> rX,UI5	Compare Immediate	<a href="#">on page 749</a>
<b>cmpl</b> crD,L,rA,rB <b>se_cmpl</b> rX,rY	Compare Logical	<a href="#">on page 752</a>
<b>e_cmpli</b> crD,rA,SCI8 <b>e_cmpl16i</b> rA,UI16 <b>se_cmpli</b> rX,UI5	Compare Logical Immediate	<a href="#">on page 752</a>

### 13.4.6 Integer select instruction

The **isel** instruction provides a means to select one of two registers and place the result in a destination register under the control of a predicate value supplied by a CR bit.

The integer select instruction is listed in [Table 247](#).

**Table 247. Integer select instruction set index**

Mnemonic	Instruction	Reference
<b>isel</b> rD,rA,rB,crb	Integer Select	EIS

### 13.4.7 Integer trap instructions

Trap instructions test for a specified set of conditions by comparing the contents of one GPR with a second GPR. If any of the conditions tested by a Trap instruction are met, a trap exception type program interrupt is invoked. If none of the tested conditions are met, instruction execution continues normally.

The contents of GPR(rA) are compared with the contents of GPR(rB). For **twi** and **tw**, only the contents of bits 32–63 of rA (and rB) participate in the comparison.

This comparison results in five conditions that are ANDed with TO. If the result is not 0, the trap exception type program interrupt is invoked. These conditions are as shown in [Table 248](#).

**Table 248. Integer trap conditions**

TO Bit	ANDed with condition
0	Less Than, using signed comparison
1	Greater Than, using signed comparison
2	Equal
3	Less Than, using unsigned comparison
4	Greater Than, using unsigned comparison

The integer trap instruction is listed in [Table 249](#).

**Table 249. Integer trap instruction set index**

Mnemonic	Instruction	Reference
<b>tw</b> TO,rA,rB	Trap Word	Book E

### 13.4.8 Integer rotate and shift instructions

Instructions are provided that perform shifts and rotates on data from a GPR and return the result, or a portion of the result, to a GPR.

The rotation operations rotate a 32-bit quantity left by a specified number of bit positions. Bits that exit from position 32 enter at position 63.

The rotate<sub>32</sub> operation is used to rotate a given 32-bit quantity.

Some rotate and shift instructions employ a mask generator. The mask is 32 bits long, and consists of 1 bits from a start bit, *mstart*, through and including a stop bit, *mstop*, and 0-bits elsewhere. The values of *mstart* and *mstop* range from 32 to 63. If *mstart* > *mstop*, the 1 bits wrap around from position 63 to position 0. Thus the mask is formed as follows:

```

if mstart ≤ mstop then
    maskmstart:mstop = ones
    maskall other bits = zeros
else
    maskmstart:63 = ones
  
```

mask<sub>32:mstop</sub> = ones  
 mask<sub>all other bits</sub> = zeros

There is no way to specify an all-zero mask.

For instructions that use the rotate<sub>32</sub> operation, the mask start and stop positions are always in bits 32–63 of the mask.

The use of the mask is described in following sections.

The rotate word and shift word instructions with Rc = 1 set the first three bits of CR field 0 as described in Book E. Rotate and shift instructions do not change the OV and SO bits. Rotate and shift instructions, except algebraic right shifts, do not change the CA bit.

The instructions in [Table 250](#) rotate the contents of a register. Depending on the instruction type, the amount of the rotation is either specified as an immediate, or contained in a GPR.

**Table 250. Integer rotate instruction set index**

Mnemonic	Instruction	Reference
<b>e_rlw</b> rA,rS,rB	Rotate Left Word	<a href="#">on page 768</a>
<b>e_rlwi</b> rA,rS,SH	Rotate Left Word Immediate	<a href="#">on page 768</a>

The instructions in [Table 251](#) rotate the contents of a register. Depending on the instruction type, the result of the rotation is either inserted into the target register under control of a mask (if a mask bit is 1, the associated bit of the rotated data is placed into the target register; if a mask bit is 0, the associated bit in the target register remains unchanged) or ANDed with a mask before being placed into the target register.

The rotate left instructions allow right-rotation of the contents of a register to be performed (in concept) by a left-rotation of 32-*n*, where *n* is the number of bits by which to rotate right. They allow right-rotation of the contents of bits 32–63 of a register to be performed (in concept) by a left-rotation of 32-*n*, where *n* is the number of bits by which to rotate right.

**Table 251. Integer rotate with mask instruction set index**

Mnemonic	Instruction	Reference
<b>e_rlwimi</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then Mask Insert	<a href="#">on page 768</a>
<b>e_rlwinm</b> rA,rS,SH,MB,ME	Rotate Left Word Immediate then AND with Mask	<a href="#">on page 769</a>

The integer shift instructions are listed in [Table 252](#).

**Table 252. Integer shift instruction set index**

Mnemonic	Instruction	Reference
<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB <b>se_slw</b> rX,rY	Shift Left Word	<a href="#">on page 770</a>
<b>e_slwi</b> rA,rS,SH <b>se_slwi</b> rX,UI5	Shift Left Word Immediate	<a href="#">on page 770</a>



Table 252. Integer shift instruction set index (continued)

Mnemonic	Instruction	Reference
<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB <b>se_sraw</b> rX,rY	Shift Right Algebraic Word	<a href="#">on page 771</a>
<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH <b>se_srawi</b> rX,UI5	Shift Right Algebraic Word Immediate	<a href="#">on page 771</a>
<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB <b>se_srw</b> rX,rY	Shift Right Word	<a href="#">on page 772</a>
<b>e_srw</b> rA,rS,SH <b>se_srw</b> rX,UI5	Shift Right Word Immediate	<a href="#">on page 772</a>

## 13.5 Storage control instructions

This section lists storage control instructions, which include the following:

- [Section 13.5.1: Storage synchronization instructions](#)
- [Section 13.5.2: Cache management instructions](#)
- [Section 13.5.3: TLB management instructions](#)

### 13.5.1 Storage synchronization instructions

The memory synchronization instructions implemented by the VLE extension are identical to those defined in Book E.

The storage synchronization instructions are listed in [Table 253](#).

Table 253. Storage synchronization instruction set index

Mnemonic	Instruction	Reference
<b>mbar</b>	Memory Barrier	Book E
<b>msync</b>	Memory Synchronize	Book E

### 13.5.2 Cache management instructions

Cache management instructions implemented by the VLE extension are identical to those defined in Book E.

The cache management instructions are listed in [Table 254](#).

Table 254. Cache management instruction set index

Mnemonic	Instruction	Reference
<b>dcba</b> rA,rB	Data Cache Block Allocate	Book E
<b>dcbf</b> rA,rB	Data Cache Block Flush	Book E
<b>dcbi</b> rA,rB	Data Cache Block Invalidate	Book E

Table 254. Cache management instruction set index

Mnemonic	Instruction	Reference
<b>dcbst</b> rA,rB	Data Cache Block Store	Book E
<b>dcbt</b> CT,rA,rB	Data Cache Block Touch	Book E
<b>dcbtIs</b> CT,rA,rB	Data Cache Block Touch and Lock Set	Book E
<b>dcbtst</b> CT,rA,rB	Data Cache Block Touch for Store	Book E
<b>dcbz</b> rA,rB	Data Cache Block set to Zero	Book E
<b>icbi</b> rA,rB	Instruction Cache Block Invalidate	Book E
<b>icbt</b> CT,rA,rB	Instruction Cache Block Touch	Book E

### 13.5.3 TLB management instructions

The TLB management instructions implemented by the VLE extension are identical to those defined in Book E and in the EIS. The TLB management instructions are listed in [Table 255](#).

Table 255. TLB management instruction set index

Mnemonic	Instruction	Reference
<b>tlbivax</b> rA,rB	TLB Invalidate Virtual Address Indexed	Book E
<b>tlbre</b>	TLB Read Entry	Book E
<b>tlbsx</b> rA,rB	TLB Search Indexed	Book E
<b>tlbsync</b>	TLB Synchronize	Book E
<b>tlbwe</b>	TLB Write Entry	Book E

### 13.5.4 Instruction alignment and byte ordering

To be recognized by the instruction decoder, an instruction fetched from memory must be placed in the pipeline with its bytes in the proper order. Book E allows instructions to be placed into memory marked as either big- or little-endian. This is manageable because Book E instructions are always word-size aligned on word boundaries. VLE instructions can be either half word or word size, and are aligned on half-word boundaries. Because of this, only big-endian instruction memory is supported when executing from a page of VLE instructions. Attempts to execute VLE instructions from a page marked as little-endian generate an instruction storage interrupt byte-ordering exception.

## 13.6 Instruction listings

This section lists instructions either defined or supported by the VLE extension.

[Table 256](#) lists instructions by instruction name.

Table 256. Instructions listed by name

Instruction	Mnemonic	Reference
Add	<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Book E
Add Carrying	<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Book E
Add Extended	<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Book E
AND with Complement	<b>andc[.]</b> rA,rS,rB <b>se_andc</b> rX,rY	Book E <a href="#">on page 743</a>
AND	<b>and[.]</b> rA,rS,rB <b>se_and[.]</b> rX,rY	Book E <a href="#">on page 743</a>
Compare	<b>cmp</b> crD,L,rA,rB <b>se_cmp</b> rX,rY	Book E <a href="#">on page 749</a>
Compare Logical	<b>cmpl</b> crD,L,rA,rB <b>se_cmpl</b> rX,rY	Book E <a href="#">on page 752</a>
Count Leading Zeros Word	<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Book E
Data Cache Block Allocate	<b>dcba</b> rA,rB	Book E
Data Cache Block Flush	<b>dcbf</b> rA,rB	Book E
Data Cache Block Invalidate	<b>dcbi</b> rA,rB	Book E
Data Cache Block Store	<b>dcbst</b> rA,rB	Book E
Data Cache Block Touch	<b>dcbt</b> CT,rA,rB	Book E
Data Cache Block Touch for Store	<b>dcbstst</b> CT,rA,rB	Book E
Data Cache Block set to Zero	<b>dcbz</b> rA,rB	Book E
Divide Word	<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Book E
Divide Word Unsigned	<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Book E
Equivalent	<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Book E

Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Extend Sign Byte	<b>extsb</b> rA,rS <b>extsb.</b> rA,rS <b>se_extsb</b> rX	Book E Book E <a href="#">on page 755</a>
Extend Sign Halfword	<b>extsh</b> rA,rS <b>extsh.</b> rA,rS <b>se_extsh</b> rX	Book E Book E <a href="#">on page 755</a>
Add Immediate Shifted	<b>e_add2is</b> rD,SI	<a href="#">on page 741</a>
Add Immediate	<b>e_addi</b> rD,rA,SCI8 <b>e_addi.</b> rD,rA,SCI8 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI <b>se_addi</b> rX,OIMM	<a href="#">on page 741</a>
Add Immediate Carrying	<b>e_addic</b> rD,rA,SCI8 <b>e_addic.</b> rD,rA,SCI8	<a href="#">on page 743</a>
AND Immediate Shifted	<b>e_and2is.</b> rD,UI	<a href="#">on page 743</a>
AND Immediate	<b>e_andi[.]</b> rA,rS,SCI8 <b>se_andi</b> rX,UI5 <b>e_and2i.</b> rD,UI	<a href="#">on page 743</a>
Branch Conditional Branch Conditional Branch Conditional & Link	<b>e_bc</b> BO32,BI32,BD15 <b>se_bc</b> BO16,BI16,BD8 <b>e_bcl</b> BO32,BI32,BD15	<a href="#">on page 745</a>
Branch Branch & Link	<b>e_b</b> BD24 <b>e_bl</b> BD24	<a href="#">on page 744</a>
Compare Halfword	<b>e_cmph</b> crD,rA,rB <b>se_cmph</b> rX,rY	<a href="#">on page 750</a>
Compare Halfword Immediate	<b>e_cmph16i</b> rA,SI16	<a href="#">on page 750</a>
Compare Halfword Logical	<b>e_cmphl</b> crD,rA,rB <b>se_cmphl</b> rX,rY	<a href="#">on page 751</a>
Compare Halfword Logical Immediate	<b>e_cmphl16i</b> rA,UI16	<a href="#">on page 751</a>
Compare Immediate	<b>e_cmpi</b> crD,rA,SCI8 <b>e_cmp16i</b> rA,SI16 <b>se_cmpi</b> rX,UI5	<a href="#">on page 749</a>
Compare Logical Immediate	<b>e_cmpli</b> crD,rA,SCI8 <b>e_cmpl16i</b> rA,UI16 <b>se_cmpli</b> rX,UI5	<a href="#">on page 752</a>
Condition Register AND	<b>e_crand</b> crbD,crbA,crbB	<a href="#">on page 753</a>
Condition Register AND with Complement	<b>e_crandc</b> crbD,crbA,crbB	<a href="#">on page 753</a>
Condition Register Equivalent	<b>e_creqv</b> crbD,crbA,crbB	<a href="#">on page 753</a>
Condition Register NAND	<b>e_crnand</b> crbD,crbA,crbB	<a href="#">on page 754</a>
Condition Register NOR	<b>e_crnor</b> crbD,crbA,crbB	<a href="#">on page 754</a>

Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Condition Register OR	<b>e_cror</b> crbD,crbA,crbB	<a href="#">on page 755</a>
Condition Register OR with Complement	<b>e_crorc</b> crbD,crbA,crbB	<a href="#">on page 755</a>
Condition Register XOR	<b>e_crxor</b> crbD,crbA,crbB	<a href="#">on page 755</a>
Load Byte and Zero	<b>e_lbz</b> rD,D(rA)	<a href="#">on page 757</a>
Load Byte and Zero with Update	<b>e_lbzu</b> rD,D8(rA)	
Load Byte and Zero (16-bit form)	<b>se_lbz</b> rZ,SD4(rX)	
Load Halfword Algebraic	<b>e_lha</b> rD,D(rA)	<a href="#">on page 758</a>
Load Halfword Algebraic with Update	<b>e_lhau</b> rD,D8(rA)	
Load Halfword and Zero	<b>e_lhz</b> rD,D(rA)	<a href="#">on page 759</a>
Load Halfword and Zero with Update	<b>e_lhzu</b> rD,D8(rA)	
Load Halfword and Zero (16-bit form)	<b>se_lhz</b> rZ,SD4(rX)	
Load Immediate	<b>e_li</b> rD,LI20 <b>se_li</b> rX,UI7	<a href="#">on page 760</a>
Load Immediate Shifted	<b>e_lis</b> rD,UI	<a href="#">on page 760</a>
Load Multiple Word	<b>e_lmw</b> rD,D8(rA)	<a href="#">on page 760</a>
Load Word and Zero	<b>e_lwz</b> rD,D(rA)	<a href="#">on page 761</a>
Load Word and Zero with Update	<b>e_lwzu</b> rD,D8(rA)	
Load Word and Zero (16-bit form)	<b>se_lwz</b> rZ,SD4(rX)	
Move Condition Register Field	<b>e_mcrf</b> crD,crS	<a href="#">on page 761</a>
Multiply Low Immediate	<b>e_mulli</b> rD,rA,SCI8 <b>e_mull2i</b> rD,SI	<a href="#">on page 763</a>
OR Immediate Shifted	<b>e_or2is</b> rD,UI	<a href="#">on page 765</a>
OR Immediate	<b>e_ori</b> [.] rA,rS,SCI8 <b>e_or2i</b> rD,UI	<a href="#">on page 765</a>
Rotate Left Word	<b>e_rlwl</b> rA,rS,rB	<a href="#">on page 768</a>
Rotate Left Word Immediate	<b>e_rlwi</b> rA,rS,SH	<a href="#">on page 768</a>
Rotate Left Word Immediate then Mask Insert	<b>e_rlwimi</b> rA,rS,SH,MB,ME	<a href="#">on page 768</a>
Rotate Left Word Immediate then AND with Mask	<b>e_rlwinm</b> rA,rS,SH,MB,ME	<a href="#">on page 769</a>
Shift Left Word Immediate	<b>e_slwl</b> rA,rS,SH <b>se_slwl</b> rX,UI5	<a href="#">on page 770</a>
Shift Right Word Immediate	<b>e_srwl</b> rA,rS,SH <b>se_srwl</b> rX,UI5	<a href="#">on page 772</a>
Store Byte	<b>e_stb</b> rS,D(rA)	<a href="#">on page 773</a>
Store Byte with Update	<b>e_stbu</b> rS,D8(rA)	
Store Byte (16-bit form)	<b>se_stb</b> rZ,SD4(rX)	

Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Store Halfword	<b>e_sth</b> rS,D(rA)	<i>on page 773</i>
Store Halfword with Update	<b>e_sthu</b> rS,D8(rA)	
Store Halfword (16-bit form)	<b>se_sth</b> rZ,SD4(rX)	
Store Multiple Word	<b>e_stmw</b> rS,D8(rA)	<i>on page 774</i>
Store Word	<b>e_stw</b> rS,D(rA)	<i>on page 775</i>
Store Word with Update	<b>e_stwu</b> rS,D8(rA)	
Store Word (16-bit form)	<b>se_stw</b> rZ,SD4(rX)	
Subtract From Immediate Carrying	<b>e_subfic</b> rD,rA,SCI8 <b>e_subfic.</b> rD,rA,SCI8	<i>on page 776</i>
XOR Immediate	<b>e_xori</b> [.] rA,rS,SCI8	<i>on page 765</i>
Instruction Cache Block Invalidate	<b>icbi</b> rA,rB	Book E
Instruction Cache Block Touch	<b>icbt</b> CT,rA,rB	Book E
Integer Select	<b>isel</b> rD,rA,rB,crb	EIS
Load Byte and Zero Indexed	<b>lbzx</b> rD,rA,rB	Book E
Load Byte and Zero with Update Indexed	<b>lbzux</b> rD,rA,rB	
Load Halfword Algebraic Indexed	<b>lhax</b> rD,rA,rB	Book E
Load Halfword Algebraic with Update Indexed	<b>lhaux</b> rD,rA,rB	
Load Halfword Byte-Reverse Indexed	<b>lhbrx</b> rD,rA,rB	Book E
Load Halfword and Zero Indexed	<b>lhzx</b> rD,rA,rB	Book E
Load Halfword and Zero with Update Indexed	<b>lhzux</b> rD,rA,rB	
Load Word And Reserve Indexed	<b>lwarx</b> rD,rA,rB	Book E
Load Word Byte-Reverse Indexed	<b>lwbrx</b> rD,rA,rB	Book E
Load Word and Zero Indexed	<b>lwzx</b> rD,rA,rB	Book E
Load Word and Zero with Update Indexed	<b>lwzux</b> rD,rA,rB	
Memory Barrier	<b>mbar</b>	Book E
Move to Condition Register from Integer Exception Register	<b>mcrxr</b> crD	Book E
Move From condition register	<b>mfcrr</b> rD	Book E
Move From Device Control Register	<b>mfdcr</b> rD,DCRN	Book E
Move From Machine State Register	<b>mfmsr</b> rD	Book E
Move From Special Purpose Register	<b>mfspir</b> rD,SPRN	Book E
Memory Synchronize	<b>msync</b>	Book E
Move to Condition Register Fields	<b>mtcrf</b> FXM,rS	Book E
Move To Device Control Register	<b>mtdcr</b> DCRN,rS	Book E
Move To Machine State Register	<b>mtmsr</b> rS	Book E
Move To Special Purpose Register	<b>mtspr</b> SPRN,rS	Book E

Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Multiply High Word	<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Book E
Multiply High Word Unsigned	<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Book E
Multiply Low Word	<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Book E
NAND	<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	Book E
Negate	<b>neg</b> rD,rA <b>se_neg</b> rX <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Book E <a href="#">on page 764</a> Book E Book E Book E
NOR	<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	Book E
OR	<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB <b>se_or</b> rX,rY	Book E Book E <a href="#">on page 765</a>
OR with Complement	<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	Book E
Add	<b>se_add</b> rX,rY	<a href="#">on page 741</a>
Bit Clear	<b>se_bclr</b> rX,UI5	<a href="#">on page 746</a>
Branch to Count Register Branch to Count Register & Link	<b>se_bctr</b> <b>se_bctrl</b>	<a href="#">on page 746</a>
Bit Generate	<b>se_bgeni</b> rX,UI5	<a href="#">on page 747</a>
Branch to Link Register Branch to Link Register & Link	<b>se_blr</b> <b>se_blrl</b>	<a href="#">on page 747</a>
Bit Mask Generate	<b>se_bmski</b> rX,UI5	<a href="#">on page 748</a>
Bit Set	<b>se_bseti</b> rX,UI5	<a href="#">on page 748</a>
Branch Branch & Link	<b>se_b</b> BD8 <b>se_bl</b> BD8	<a href="#">on page 744</a>
Bit Test Immediate	<b>se_btsti</b> rX,UI5	<a href="#">on page 748</a>
Extend with Zeros Byte	<b>se_extzb</b> rX	<a href="#">on page 756</a>
Extend with Zeros Halfword	<b>se_extzh</b> rX	<a href="#">on page 756</a>
Instruction Synchronize	<b>se_isync</b>	<a href="#">on page 757</a>
Move from Alternate Register	<b>se_mfar</b> rX,arY	<a href="#">on page 762</a>
Move From Count Register	<b>se_mfctr</b> rX	<a href="#">on page 762</a>

Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Move From Link Register	<b>se_mflr</b> rX	<a href="#">on page 762</a>
Move Register	<b>se_mr</b> rX,rY	<a href="#">on page 762</a>
Move to Alternate Register	<b>se_mtar</b> arX,rY	<a href="#">on page 763</a>
Move To Count Register	<b>se_mtctr</b> rX	<a href="#">on page 763</a>
Move To Link Register	<b>se_mtlr</b> rX	<a href="#">on page 763</a>
Multiply Low Word	<b>se_mullw</b> rX,rY	<a href="#">on page 764</a>
NOT	<b>se_not</b> rX	<a href="#">on page 764</a>
Subtract	<b>se_sub</b> rX,rY	<a href="#">on page 775</a>
Subtract From	<b>se_subf</b> rX,rY	<a href="#">on page 776</a>
Subtract Immediate	<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	<a href="#">on page 776</a>
Shift Left Word	<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB <b>se_slw</b> rX,rY	Book E Book E <a href="#">on page 770</a>
Shift Right Algebraic Word	<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB <b>se_sraw</b> rX,rY	Book E Book E <a href="#">on page 771</a>
Shift Right Algebraic Word Immediate	<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH <b>se_srawi</b> rX,UI5	Book E Book E <a href="#">on page 771</a>
Shift Right Word	<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB <b>se_srw</b> rX,rY	Book E Book E <a href="#">on page 772</a>
Store Byte Indexed Store Byte with Update Indexed	<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Book E
Store Halfword Byte-Reverse Indexed	<b>sthbrx</b> rS,rA,rB	Book E
Store Halfword Indexed Store Halfword with Update Indexed	<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Book E
Store Word Byte-Reverse Indexed	<b>stwbrx</b> rS,rA,rB	Book E
Store Word Conditional Indexed	<b>stwcx.</b> rS,rA,rB	Book E
Store Word Indexed Store Word with Update Indexed	<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Book E
Subtract From	<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Book E



Table 256. Instructions listed by name (continued)

Instruction	Mnemonic	Reference
Subtract From Carrying	<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Book E
TLB Invalidate Virtual Address Indexed	<b>tlbivax</b> rA,rB	Book E
TLB Read Entry	<b>tlbre</b>	Book E
TLB Search Indexed	<b>tlbsx</b> rA,rB	Book E
TLB Synchronize	<b>tlbsync</b>	Book E
TLB Write Entry	<b>tlbwe</b>	Book E
Trap Word	<b>tw</b> TO,rA,rB	Book E
Write MSR External Enable	<b>wrtee</b> rA	Book E
Write MSR External Enable Immediate	<b>wrteei</b> E	Book E
XOR	<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	Book E

[Table 257](#) lists instructions by mnemonic.

Table 257. Instructions listed by mnemonic

Mnemonic	Instruction	Reference
<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Add	Book E
<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Add Carrying	Book E
<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Add Extended	Book E
<b>andc[.]</b> rA,rS,rB	AND with Complement	Book E
<b>and[.]</b> rA,rS,rB	AND	Book E
<b>cmp</b> crD,L,rA,rB	Compare	Book E
<b>cmpl</b> crD,L,rA,rB	Compare Logical	Book E
<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Count Leading Zeros Word	Book E
<b>dcba</b> rA,rB	Data Cache Block Allocate	Book E
<b>dcbf</b> rA,rB	Data Cache Block Flush	Book E

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>dcbi</b> rA,rB	Data Cache Block Invalidate	Book E
<b>dcbst</b> rA,rB	Data Cache Block Store	Book E
<b>dcbt</b> CT,rA,rB	Data Cache Block Touch	Book E
<b>dcbtst</b> CT,rA,rB	Data Cache Block Touch for Store	Book E
<b>dcbz</b> rA,rB	Data Cache Block set to Zero	Book E
<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Divide Word	Book E
<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Divide Word Unsigned	Book E
<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Equivalent	Book E
<b>extsb</b> rA,rS <b>extsb.</b> rA,rS	Extend Sign Byte	Book E
<b>extsh</b> rA,rS <b>extsh.</b> rA,rS	Extend Sign Halfword	Book E
<b>e_add2is</b> rD,SI	Add Immediate Shifted	<a href="#">on page 741</a>
<b>e_addi</b> rD,rA,SCI8 <b>e_addi.</b> rD,rA,SCI8 <b>e_add16i</b> rD,rA,SI <b>e_add2i.</b> rD,SI	Add Immediate	<a href="#">on page 741</a>
<b>e_addic</b> rD,rA,SCI8 <b>e_addic.</b> rD,rA,SCI8	Add Immediate Carrying	<a href="#">on page 743</a>
<b>e_and2is.</b> rD,UI	AND Immediate Shifted	<a href="#">on page 743</a>
<b>e_andi[.]</b> rA,rS,SCI8 <b>e_and2i.</b> rD,UI	AND Immediate	<a href="#">on page 743</a>
<b>e_bc</b> BO32,BI32,BD15 <b>e_bcl</b> BO32,BI32,BD15	Branch Conditional Branch Conditional & Link	<a href="#">on page 745</a>
<b>e_b</b> BD24 <b>e_bl</b> BD24	Branch Branch & Link	<a href="#">on page 744</a>
<b>e_cmph</b> crD,rA,rB	Compare Halfword	<a href="#">on page 750</a>
<b>e_cmph16i</b> rA,SI16	Compare Halfword Immediate	<a href="#">on page 750</a>
<b>e_cmphl</b> crD,rA,rB	Compare Halfword Logical	<a href="#">on page 751</a>
<b>e_cmphl16i</b> rA,UI16	Compare Halfword Logical Immediate	<a href="#">on page 751</a>
<b>e_cmpi</b> crD,rA,SCI8 <b>e_cmp16i</b> rA,SI16	Compare Immediate	<a href="#">on page 749</a>

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
e_cmpli crD,rA,SCI8 e_cmpl16i rA,UI16	Compare Logical Immediate	<a href="#">on page 752</a>
e_crand crbD,crbA,crbB	Condition Register AND	<a href="#">on page 753</a>
e_crandc crbD,crbA,crbB	Condition Register AND with Complement	<a href="#">on page 753</a>
e_creqv crbD,crbA,crbB	Condition Register Equivalent	<a href="#">on page 753</a>
e_crnand crbD,crbA,crbB	Condition Register NAND	<a href="#">on page 754</a>
e_crnor crbD,crbA,crbB	Condition Register NOR	<a href="#">on page 754</a>
e_cror crbD,crbA,crbB	Condition Register OR	<a href="#">on page 755</a>
e_crorc crbD,crbA,crbB	Condition Register OR with Complement	<a href="#">on page 755</a>
e_crxor crbD,crbA,crbB	Condition Register XOR	<a href="#">on page 755</a>
e_lbz rD,D(rA) e_lbzu rD,D8(rA)	Load Byte and Zero Load Byte and Zero with Update	<a href="#">on page 757</a>
e_lha rD,D(rA) e_lhau rD,D8(rA)	Load Halfword Algebraic Load Halfword Algebraic with Update	<a href="#">on page 758</a>
e_lhz rD,D(rA) e_lhzu rD,D8(rA)	Load Halfword and Zero Load Halfword and Zero with Update	<a href="#">on page 759</a>
e_li rD,LI20	Load Immediate	<a href="#">on page 760</a>
e_lis rD,UI	Load Immediate Shifted	<a href="#">on page 760</a>
e_lmw rD,D8(rA)	Load Multiple Word	<a href="#">on page 761</a>
e_lwz rD,D(rA) e_lwzu rD,D8(rA)	Load Word and Zero Load Word and Zero with Update	<a href="#">on page 761</a>
e_mcrf crD,crS	Move Condition Register Field	<a href="#">on page 763</a>
e_mulli rD,rA,SCI8 e_mull2i rD,SI	Multiply Low Immediate	<a href="#">on page 765</a>
e_or2is rD,UI	OR Immediate Shifted	<a href="#">on page 765</a>
e_ori[.] rA,rS,SCI8 e_or2i rD,UI	OR Immediate	<a href="#">on page 768</a>
e_rlw rA,rS,rB	Rotate Left Word	<a href="#">on page 768</a>
e_rlwi rA,rS,SH	Rotate Left Word Immediate	<a href="#">on page 768</a>
e_rlwimi rA,rS,SH,MB,ME	Rotate Left Word Immediate then Mask Insert	<a href="#">on page 769</a>
e_rlwinm rA,rS,SH,MB,ME	Rotate Left Word Immediate then AND with Mask	<a href="#">on page 770</a>
e_slwi rA,rS,SH	Shift Left Word Immediate	<a href="#">on page 761</a>
e_srwi rA,rS,SH	Shift Right Word Immediate	Book E
e_stb rS,D(rA) e_stbu rS,D8(rA)	Store Byte Store Byte with Update	<a href="#">on page 773</a>

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>e_sth</b> rS,D(rA)	Store Halfword	<a href="#">on page 773</a>
<b>e_sthu</b> rS,D8(rA)	Store Halfword with Update	<a href="#">on page 773</a>
<b>e_stmw</b> rS,D8(rA)	Store Multiple Word	<a href="#">on page 774</a>
<b>e_stw</b> rS,D(rA)	Store Word	<a href="#">on page 775</a>
<b>e_stwu</b> rS,D8(rA)	Store Word with Update	<a href="#">on page 775</a>
<b>e_subfic</b> rD,rA,SCI8 <b>e_subfic.</b> rD,rA,SCI8	Subtract From Immediate Carrying	<a href="#">on page 776</a>
<b>e_xori</b> [.] rA,rS,SCI8	XOR Immediate	<a href="#">on page 765</a>
<b>icbi</b> rA,rB	Instruction Cache Block Invalidate	Book E
<b>icbt</b> CT,rA,rB	Instruction Cache Block Touch	Book E
<b>isel</b> rD,rA,rB,crb	Integer Select	EIS
<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	Book E
<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	Book E
<b>lhbrx</b> rD,rA,rB	Load Halfword Byte-Reverse Indexed	Book E
<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	Book E
<b>lwarx</b> rD,rA,rB	Load Word And Reserve Indexed	Book E
<b>lwbrx</b> rD,rA,rB	Load Word Byte-Reverse Indexed	Book E
<b>lwzx</b> rD,rA,rB <b>lwzux</b> rD,rA,rB	Load Word and Zero Indexed Load Word and Zero with Update Indexed	Book E
<b>mbar</b>	Memory Barrier	Book E
<b>mcrxr</b> crD	Move to Condition Register from Integer Exception Register	Book E
<b>mfcrr</b> rD	Move From condition register	Book E
<b>mfdcrr</b> rD,DCRN	Move From Device Control Register	Book E
<b>mfmsrr</b> rD	Move From Machine State Register	Book E
<b>mfsprr</b> rD,SPRN	Move From Special Purpose Register	Book E
<b>msync</b>	Memory Synchronize	Book E
<b>mtcrf</b> FXM,rS	Move to Condition Register Fields	Book E
<b>mtdcrr</b> DCRN,rS	Move To Device Control Register	Book E
<b>mtmsrr</b> rS	Move To Machine State Register	Book E
<b>mtsprr</b> SPRN,rS	Move To Special Purpose Register	Book E
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word	Book E

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word Unsigned	Book E
<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Multiply Low Word	Book E
<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	NAND	Book E
<b>neg</b> rD,rA <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Negate	Book E
<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	NOR	Book E
<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB	OR	Book E
<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	OR with Complement	Book E
<b>se_add</b> rX,rY	Add	<a href="#">on page 741</a>
<b>se_addi</b> rX,OIMM	Add Immediate	<a href="#">on page 741</a>
<b>se_andc</b> rX,rY	AND with Complement	<a href="#">on page 743</a>
<b>se_andi</b> rX,UI5	AND Immediate	<a href="#">on page 743</a>
<b>se_and[.]</b> rX,rY	AND	<a href="#">on page 743</a>
<b>se_bc</b> BO16,BI16,BD8	Branch Conditional	<a href="#">on page 745</a>
<b>se_bclri</b> rX,UI5	Bit Clear	<a href="#">on page 746</a>
<b>se_bctr</b> <b>se_bctrl</b>	Branch to Count Register Branch to Count Register & Link	<a href="#">on page 746</a>
<b>se_bgeni</b> rX,UI5	Bit Generate	<a href="#">on page 746</a>
<b>se_blr</b> <b>se_blrl</b>	Branch to Link Register Branch to Link Register & Link	<a href="#">on page 747</a>
<b>se_bmski</b> rX,UI5	Bit Mask Generate	<a href="#">on page 747</a>
<b>se_bseti</b> rX,UI5	Bit Set	<a href="#">on page 748</a>
<b>se_b</b> BD8 <b>se_bl</b> BD8	Branch Branch & Link	<a href="#">on page 748</a>
<b>se_btsti</b> rX,UI5	Bit Test Immediate	<a href="#">on page 744</a>
<b>se_cmp</b> rX,rY	Compare	<a href="#">on page 749</a>
<b>se_cmph</b> rX,rY	Compare Halfword	<a href="#">on page 750</a>
<b>se_cmphi</b> rX,rY	Compare Halfword Logical	<a href="#">on page 751</a>

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>se_cmpi</b> rX,UI5	Compare Immediate	<a href="#">on page 749</a>
<b>se_cmpl</b> rX,rY	Compare Logical	<a href="#">on page 752</a>
<b>se_cmpli</b> rX,UI5	Compare Logical Immediate	<a href="#">on page 752</a>
<b>se_extsb</b> rX	Extend Sign Byte	<a href="#">on page 755</a>
<b>se_extsh</b> rX	Extend Sign Halfword	<a href="#">on page 755</a>
<b>se_extzb</b> rX	Extend with Zeros Byte	<a href="#">on page 756</a>
<b>se_extzh</b> rX	Extend with Zeros Halfword	<a href="#">on page 756</a>
<b>se_isync</b>	Instruction Synchronize	<a href="#">on page 757</a>
<b>se_lbz</b> rZ,SD4(rX)	Load Byte and Zero (16-bit form)	<a href="#">on page 757</a>
<b>se_lhz</b> rZ,SD4(rX)	Load Halfword and Zero (16-bit form)	<a href="#">on page 759</a>
<b>se_li</b> rX,UI7	Load Immediate	<a href="#">on page 760</a>
<b>se_lwz</b> rZ,SD4(rX)	Load Word and Zero (16-bit form)	<a href="#">on page 761</a>
<b>se_mfar</b> rX,arY	Move from Alternate Register	<a href="#">on page 762</a>
<b>se_mfctr</b> rX	Move From Count Register	<a href="#">on page 762</a>
<b>se_mflr</b> rX	Move From Link Register	<a href="#">on page 762</a>
<b>se_mr</b> rX,rY	Move Register	<a href="#">on page 762</a>
<b>se_mtar</b> arX,rY	Move to Alternate Register	<a href="#">on page 763</a>
<b>se_mtctr</b> rX	Move To Count Register	<a href="#">on page 763</a>
<b>se_mtlr</b> rX	Move To Link Register	<a href="#">on page 763</a>
<b>se_mullw</b> rX,rY	Multiply Low Word	<a href="#">on page 764</a>
<b>se_neg</b> rX	Negate	<a href="#">on page 764</a>
<b>se_not</b> rX	NOT	<a href="#">on page 764</a>
<b>se_or</b> rX,rY	OR	<a href="#">on page 765</a>
<b>se_slw</b> rX,rY	Shift Left Word	<a href="#">on page 770</a>
<b>se_slwi</b> rX,UI5	Shift Left Word Immediate	<a href="#">on page 770</a>
<b>se_sraw</b> rX,rY	Shift Right Algebraic Word	<a href="#">on page 771</a>
<b>se_srawi</b> rX,UI5	Shift Right Algebraic Word Immediate	<a href="#">on page 771</a>
<b>se_srw</b> rX,rY	Shift Right Word	<a href="#">on page 772</a>
<b>se_srwi</b> rX,UI5	Shift Right Word Immediate	<a href="#">on page 772</a>
<b>se_stb</b> rZ,SD4(rX)	Store Byte (16-bit form)	<a href="#">on page 773</a>
<b>se_sth</b> rZ,SD4(rX)	Store Halfword (16-bit form)	<a href="#">on page 773</a>
<b>se_stw</b> rZ,SD4(rX)	Store Word (16-bit form)	<a href="#">on page 775</a>
<b>se_sub</b> rX,rY	Subtract	<a href="#">on page 775</a>
<b>se_subf</b> rX,rY	Subtract From	<a href="#">on page 776</a>

Table 257. Instructions listed by mnemonic (continued)

Mnemonic	Instruction	Reference
<b>se_subi</b> rX,OIMM <b>se_subi.</b> rX,OIMM	Subtract Immediate	<a href="#">on page 776</a>
<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB	Shift Left Word	Book E
<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB	Shift Right Algebraic Word	Book E
<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH	Shift Right Algebraic Word Immediate	Book E
<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB	Shift Right Word	Book E
<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Store Byte Indexed Store Byte with Update Indexed	Book E
<b>sthbrx</b> rS,rA,rB	Store Halfword Byte-Reverse Indexed	Book E
<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Store Halfword Indexed Store Halfword with Update Indexed	Book E
<b>stwbrx</b> rS,rA,rB	Store Word Byte-Reverse Indexed	Book E
<b>stwcx.</b> rS,rA,rB	Store Word Conditional Indexed	Book E
<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Store Word Indexed Store Word with Update Indexed	Book E
<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Subtract From	Book E
<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Subtract From Carrying	Book E
<b>tlbivax</b> rA,rB	TLB Invalidate Virtual Address Indexed	Book E
<b>tlbre</b>	TLB Read Entry	Book E
<b>tlbsx</b> rA,rB	TLB Search Indexed	Book E
<b>tlbsync</b>	TLB Synchronize	Book E
<b>tlbwe</b>	TLB Write Entry	Book E
<b>tw</b> TO,rA,rB	Trap Word	Book E
<b>wrtee</b> rA	Write MSR External Enable	Book E
<b>wrteei</b> E	Write MSR External Enable Immediate	Book E
<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	XOR	Book E

## 14 VLE instruction set

The VLE extension ISA is defined in the instruction pages in this chapter. Because of the various immediate field and displacement field calculations used in the VLE extension, a description of the less obvious ones precedes the actual instruction pages, and the instruction descriptions generally assume the appropriate calculation has been performed.

*Note:* The instructions in this section are listed in order of the root instruction. For example, **e\_cmpi** and **se\_cmpi** are both listed under **cmpi**.

### 14.1 Book E– and EIS-defined instructions

[Table 258](#) lists instructions that are used by the VLE extension that are defined by Book E or the EIS. Full descriptions of those instructions can be found in the EREF.

Descriptions in this chapter indicate any limitations on the behavior of VLE instructions as compared to their Book E and EIS equivalents.

**Table 258. Book E– and EIS-defined instructions listed by mnemonic**

Mnemonic	Instruction	Defining architecture
<b>add</b> rD,rA,rB <b>add.</b> rD,rA,rB <b>addo</b> rD,rA,rB <b>addo.</b> rD,rA,rB	Add	Book E
<b>addc</b> rD,rA,rB <b>addc.</b> rD,rA,rB <b>addco</b> rD,rA,rB <b>addco.</b> rD,rA,rB	Add Carrying	Book E
<b>adde</b> rD,rA,rB <b>adde.</b> rD,rA,rB <b>addeo</b> rD,rA,rB <b>addeo.</b> rD,rA,rB	Add Extended	Book E
<b>andc</b> [.] rA,rS,rB	AND with Complement	Book E
<b>and</b> [.] rA,rS,rB	AND	Book E
<b>cmp</b> crD,L,rA,rB	Compare	Book E
<b>cmpl</b> crD,L,rA,rB	Compare Logical	Book E
<b>cntlzw</b> rA,rS <b>cntlzw.</b> rA,rS	Count Leading Zeros Word	Book E
<b>dcba</b> rA,rB	Data Cache Block Allocate	Book E
<b>dcbf</b> rA,rB	Data Cache Block Flush	Book E
<b>dcbi</b> rA,rB	Data Cache Block Invalidate	Book E
<b>dcbst</b> rA,rB	Data Cache Block Store	Book E
<b>dcbt</b> CT,rA,rB	Data Cache Block Touch	Book E



Table 258. Book E– and EIS-defined instructions listed by mnemonic (continued)

Mnemonic	Instruction	Defining architecture
<b>dcbtls</b> CT,rA,rB	Data Cache Block Touch and Lock Set	Book E
<b>dcbtst</b> CT,rA,rB	Data Cache Block Touch for Store	Book E
<b>dcbz</b> rA,rB	Data Cache Block set to Zero	Book E
<b>divw</b> rD,rA,rB <b>divw.</b> rD,rA,rB <b>divwo</b> rD,rA,rB <b>divwo.</b> rD,rA,rB	Divide Word	Book E
<b>divwu</b> rD,rA,rB <b>divwu.</b> rD,rA,rB <b>divwuo</b> rD,rA,rB <b>divwuo.</b> rD,rA,rB	Divide Word Unsigned	Book E
<b>eqv</b> rA,rS,rB <b>eqv.</b> rA,rS,rB	Equivalent	Book E
<b>extsb</b> rA,rS <b>extsb.</b> rA,rS	Extend Sign Byte	Book E
<b>extsh</b> rA,rS <b>extsh.</b> rA,rS	Extend Sign Halfword	Book E
<b>e_srwi</b> rA,rS,SH	Shift Right Word Immediate	Book E
<b>icbi</b> rA,rB	Instruction Cache Block Invalidate	Book E
<b>icbt</b> CT,rA,rB	Instruction Cache Block Touch	Book E
<b>lbzx</b> rD,rA,rB <b>lbzux</b> rD,rA,rB	Load Byte and Zero Indexed Load Byte and Zero with Update Indexed	Book E
<b>lhax</b> rD,rA,rB <b>lhaux</b> rD,rA,rB	Load Halfword Algebraic Indexed Load Halfword Algebraic with Update Indexed	Book E
<b>lhrx</b> rD,rA,rB	Load Halfword Byte-Reverse Indexed	Book E
<b>lhzx</b> rD,rA,rB <b>lhzux</b> rD,rA,rB	Load Halfword and Zero Indexed Load Halfword and Zero with Update Indexed	Book E
<b>lwarx</b> rD,rA,rB	Load Word And Reserve Indexed	Book E
<b>lwbrx</b> rD,rA,rB	Load Word Byte-Reverse Indexed	Book E
<b>lwzx</b> rD,rA,rB <b>lwzux</b> rD,rA,rB	Load Word and Zero Indexed Load Word and Zero with Update Indexed	Book E
<b>mbar</b>	Memory Barrier	Book E
<b>mcrxr</b> crD	Move to Condition Register from Integer Exception Register	Book E
<b>mfcrr</b> rD	Move From condition register	Book E
<b>mfocr</b> rD,DCRN	Move From Device Control Register	Book E
<b>mfmsr</b> rD	Move From Machine State Register	Book E
<b>mfpsr</b> rD,SPRN	Move From Special Purpose Register	Book E

Table 258. Book E– and EIS-defined instructions listed by mnemonic (continued)

Mnemonic	Instruction	Defining architecture
<b>msync</b>	Memory Synchronize	Book E
<b>mtcrf</b> FXM,rS	Move to Condition Register Fields	Book E
<b>mtdcr</b> DCRN,rS	Move To Device Control Register	Book E
<b>mtmsr</b> rS	Move To Machine State Register	Book E
<b>mtspr</b> SPRN,rS	Move To Special Purpose Register	Book E
<b>mulhw</b> rD,rA,rB <b>mulhw.</b> rD,rA,rB	Multiply High Word	Book E
<b>mulhwu</b> rD,rA,rB <b>mulhwu.</b> rD,rA,rB	Multiply High Word Unsigned	Book E
<b>mullw</b> rD,rA,rB <b>mullw.</b> rD,rA,rB <b>mullwo</b> rD,rA,rB <b>mullwo.</b> rD,rA,rB	Multiply Low Word	Book E
<b>nand</b> rA,rS,rB <b>nand.</b> rA,rS,rB	NAND	Book E
<b>neg</b> rD,rA <b>neg.</b> rD,rA <b>nego</b> rD,rA <b>nego.</b> rD,rA	Negate	Book E
<b>nor</b> rA,rS,rB <b>nor.</b> rA,rS,rB	NOR	Book E
<b>or</b> rA,rS,rB <b>or.</b> rA,rS,rB	OR	Book E
<b>orc</b> rA,rS,rB <b>orc.</b> rA,rS,rB	OR with Complement	Book E
<b>slw</b> rA,rS,rB <b>slw.</b> rA,rS,rB	Shift Left Word	Book E
<b>sraw</b> rA,rS,rB <b>sraw.</b> rA,rS,rB	Shift Right Algebraic Word	Book E
<b>srawi</b> rA,rS,SH <b>srawi.</b> rA,rS,SH	Shift Right Algebraic Word Immediate	Book E
<b>srw</b> rA,rS,rB <b>srw.</b> rA,rS,rB	Shift Right Word	Book E
<b>stbx</b> rS,rA,rB <b>stbux</b> rS,rA,rB	Store Byte Indexed Store Byte with Update Indexed	Book E
<b>sthbrx</b> rS,rA,rB	Store Halfword Byte-Reverse Indexed	Book E
<b>sthx</b> rS,rA,rB <b>sthux</b> rS,rA,rB	Store Halfword Indexed Store Halfword with Update Indexed	Book E

**Table 258. Book E– and EIS-defined instructions listed by mnemonic (continued)**

<b>Mnemonic</b>	<b>Instruction</b>	<b>Defining architecture</b>
<b>stwbrx</b> rS,rA,rB	Store Word Byte-Reverse Indexed	Book E
<b>stwcx.</b> rS,rA,rB	Store Word Conditional Indexed	Book E
<b>stwx</b> rS,rA,rB <b>stwux</b> rS,rA,rB	Store Word Indexed Store Word with Update Indexed	Book E
<b>subf</b> rD,rA,rB <b>subf.</b> rD,rA,rB <b>subfo</b> rD,rA,rB <b>subfo.</b> rD,rA,rB	Subtract From	Book E
<b>subfc</b> rD,rA,rB <b>subfc.</b> rD,rA,rB <b>subfco</b> rD,rA,rB <b>subfco.</b> rD,rA,rB	Subtract From Carrying	Book E
<b>tlbivax</b> rA,rB	TLB Invalidate Virtual Address Indexed	Book E
<b>tlbre</b>	TLB Read Entry	Book E
<b>tlbsx</b> rA,rB	TLB Search Indexed	Book E
<b>tlbsync</b>	TLB Synchronize	Book E
<b>tlbwe</b>	TLB Write Entry	Book E
<b>tw</b> TO,rA,rB	Trap Word	Book E
<b>wrtee</b> rA	Write MSR External Enable	Book E
<b>wrteei</b> E	Write MSR External Enable Immediate	Book E
<b>xor</b> rA,rS,rB <b>xor.</b> rA,rS,rB	XOR	Book E
<b>isel</b> rD,rA,rB,crb	Integer Select	EIS

## 14.2 Immediate field and displacement field encodings

Table 259 shows encodings for immediate and displacement fields.

**Table 259. Immediate field and displacement field encodings**

Encoding	Description
BD15	Format used by 32-bit branch conditional class instructions. The BD15 field is an 15-bit signed displacement which is sign-extended and shifted left one bit (concatenated with 0b0) and then added to the current instruction address to form the branch target address.
BD24	Format used by 32-bit branch class instructions. The BD24 field is an 24-bit signed displacement which is sign-extended and shifted left one bit (concatenated with 0b0) and then added to the current instruction address to form the branch target address.
BD8	Format used by 16-bit branch and branch conditional class instructions. The BD8 field is an 8-bit signed displacement which is sign-extended and shifted left one bit (concatenated with 0b0) and then added to the current instruction address to form the branch target address.
D	Format used by some 32-bit load and store class instructions. The D field is a 16-bit signed displacement which is sign-extended to 32 bits, and then added to the base register to form a 32-bit EA.
D8	Format used by some 32-bit load and store class instructions. The D8 field is a 8-bit signed displacement which is sign-extended to 32 bits, and then added to the base register to form a 32-bit EA.
F, SCL, UI8 (SCI8 format)	Format used by some 32-bit arithmetic, compare, and logical instructions. The UI8 field is an 8-bit immediate value shifted left 0, 1, 2, or 3 byte positions according to the value of the SCL field. The remaining bits in the 32-bit word are filled with the value of the F field, and the resulting 32-bit value is used as one operand of the instruction. More formally, if SCL=0 then imm_value ← $^{24}F \parallel UI8$ else if SCL=1 then imm_value ← $^{16}F \parallel UI8 \parallel ^8F$ else if SCL=2 then imm_value ← $^8F \parallel UI8 \parallel ^{16}F$ else imm_value ← $UI8 \parallel ^{24}F$
LI20	Format used by 32-bit <b>e_li</b> instruction. The LI20 field is a 20-bit signed displacement which is sign-extended to 32 bits for the <b>e_li</b> instruction.
OIM5	Format used by the 16-bit <b>se_addi</b> , <b>se_cmpli</b> , and <b>se_subi[.]</b> instructions. The OIM5 instruction field is a 5-bit value in the range 0–31 and is used to represent immediate values in the range 1–32, thus the binary encoding of 0b00000 represents an immediate value of 1, 0b00001 represents an immediate value of 2, and so on. In the instruction descriptions, OIMM represents the immediate value, not the OIM5 instruction field binary encoding.
SCI8 format	Refer to F, SCL, UI8 (SCI8 format)
SD4	Format used by 16-bit load and store class instructions. The SD4 field is a 4-bit unsigned immediate value zero-extended to 32 bits, shifted left according to the size of the operation, and then added to the base register to form a 32-bit EA. For byte operations, no shift is performed. For half-word operations, the immediate is shifted left one bit (concatenated with 0b0). For word operations, the immediate is shifted left two bits (concatenated with 0b00). For future double-word operations, the immediate is shifted left three bits (concatenated with 0b000).
SI (D format, I16A format)	Format used by certain 32-bit arithmetic type instructions. The SI field is a 16-bit signed immediate value sign-extended to 32 bits and used as one operand of the instruction. The instruction encoding differs between the D and I16A instruction formats

**Table 259. Immediate field and displacement field encodings (continued)**


Encoding	Description
UI (I16A, I16L formats)	Format used by certain 32-bit logical and arithmetic type instructions. The UI field is a 16-bit unsigned immediate value zero-extended to 32 bits or padded with 16 zeros and used as one operand of the instruction. The instruction encoding differs between the I16A and I16L instruction formats.
UI5	This format is used by some 16-bit Reg+Imm class instructions. The UI5 field is a 5-bit unsigned immediate value zero-extended to 32 bits and used as the second operand of the instruction. For other 16-bit Reg+Imm class instructions, the UI5 field is a 5-bit unsigned immediate value used to select a register bit in the range 0–31.
UI7	This format is used by the 16-bit <b>se_li</b> instructions. The UI7 field is a 7-bit unsigned immediate value zero-extended to 32 bits and used as the operand of the instruction.

**\_addx** **\_addx**

VLE	User
-----	------

**Add**

**se\_add** **rX,rY**



$$\text{sum}_{32:63} \leftarrow \text{GPR}(\text{RX}) + \text{GPR}(\text{RY})$$

$$\text{GPR}(\text{RX}) \leftarrow \text{sum}_{32:63}$$

The sum of the contents of GPR(rX) and the contents of GPR(rY) is placed into GPR(rX).


Special registers altered: None

**\_addix** **\_addix**

VLE	User
-----	------

Add [2 operand] Immediate [Shifted] [and Record]

**e\_add16i** **rD,rA,SI**



$$a \leftarrow \text{GPR}(\text{RA})$$


$$b \leftarrow \text{EXTS}(\text{SI})$$

$$\text{GPR}(\text{RD}) \leftarrow a + b$$

The sum of the contents of GPR(rA) and the sign-extended value of field SI is placed into GPR(rD).

Special Registers Altered: None

**e\_add2i.** **rA,SI**



$$\text{SI} \leftarrow \text{SI}_{0:4} \parallel \text{SI}_{5:15}$$

$$\text{sum}_{32:63} \leftarrow \text{GPR}(\text{RA}) + \text{EXTS}(\text{SI})$$



The sum of the contents of GPR(rX) and the zero-extended offset value of OIM5 (a final value in the range 1–32), is placed into GPR(rX).

Special Registers Altered: None

**\_addicx**

VLE	User
-----	------

**\_addicx**

Add Immediate Carrying [and Record]

**e\_addic** **rD,rA,SCI8** (Rc = 0)  
**e\_addic.** **rD,rA,SCI8** (Rc = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	1	0	RD			RA			1	0	0	1	Rc	F	SCL	UI8												

```

imm ← SCI8(F,SCL,UI8)
carry32:63 ← Carry(GPR(RA) + imm)
sum32:63 ← GPR(RA) + imm
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RD) ← sum32:63
CA ← carry32
    
```

The sum of the contents of GPR(rA) and the value of SCI8 is placed into GPR(rD).

Special Registers Altered: CA, CR0 (if Rc=1)

**\_andx**

VLE	User
-----	------

**\_andx**

AND [2 operand] [Immediate | with Complement] [and Record]

**se\_and** **rX,rY** (Rc = 0)  
**se\_and.** **rX,rY** (Rc = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	1	0	0	0	1	1	Rc	RY			RX				

**e\_and2i.** **rD,UI**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	0	RD			UI <sub>0:4</sub>			1	1	0	0	1	UI <sub>5:15</sub>														

**e\_and2is.** **rD,UI**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	0	RD			UI <sub>0:4</sub>			1	1	1	0	1	UI <sub>5:15</sub>														

**e\_andi** **rA,rS,SCI8** (Rc = 0)  
**e\_andi.** **rA,rS,SCI8** (Rc = 1)



0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	1	0	RS						RA						1	1	0	0	Rc	F	SCL	UI8						

**se\_andi** rX,UI5

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	1	0	1	1	1	UI5						RX		

**se\_andc** rX,rY

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	1	0	0	0	1	0	1	RY						RX		

```

if 'e_andi[.]' then b ← SCI8(F,SCL,UI8)
if 'se_andi' then b ← UI5
if 'se_and[.]' then b ← GPR(RY)
if 'se_andc' then b ← ¬GPR(RY)
if 'e_and2i.' then b ← 160 || UI0:4 || UI5:15
if 'e_and2is.' then b ← UI0:4 || UI5:15 || 160
result32:63 ← GPR(RS or RD or RX) & b
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
if 'se_and[ci]' then GPR(RX) ← result32:63 else GPR(RA or RD) ← result32:63
    
```

For **e\_andi[.]**, the contents of GPR(rS) are ANDed with the value of SCI8.  
 For **e\_and2i.**, the contents of GPR(rD) are ANDed with 160 || UI.  
 For **e\_and2is.**, the contents of GPR(rD) are ANDed with UI || 160.  
 For **se\_andi**, the contents of GPR(rX) are ANDed with the value of UI5.  
 For **se\_and[.]**, the contents of GPR(rX) are ANDed with the contents of GPR(rY).  
 For **se\_andc**, the contents of GPR(rX) are ANDed with the one's complement of the contents of GPR(rY).  
 The result is placed into GPR(rA) or GPR(rX) (**se\_and[ic][.]**)  
 Special Registers Altered: CR0 (if Rc = 1)

<b>_bx</b>	VLE	User	<b>_bx</b>
Branch [and Link]			
<b>e_b</b>	BD24		(LK = 0)
<b>e_bl</b>	BD24		(LK = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	0	0	BD24														LK										

```

a ← CIA
NIA ← (a + EXTS(BD24||0b0))32:63
    
```



if LK=1 then LR ← CIA + 4

Let the BTEA be calculated as follows:


- For **e\_b[l]**, let BTEA be the sum of the CIA and the sign-extended value of the BD24 instruction field concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA+4 is placed into the LR.

Special Registers Altered: LR (if LK = 1)

<b>se_b</b>	BD8	(LK = 0)
<b>se_bl</b>	BD8	(LK = 1)



a ← CIA

NIA ← (a + EXTS(BD8||0b0))<sub>32:63</sub>

if LK=1 then LR ← CIA + 2

Let the BTEA be calculated as follows:

- For **se\_b[l]**, let BTEA be the sum of the CIA and the sign-extended value of the BD8 instruction field concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.


If LK = 1, the sum CIA+2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)



Branch Conditional [and Link]

<b>e_bc</b>	BO32,BI32,BD15	(LK = 0)
<b>e_bcl</b>	BO32,BI32,BD15	(LK = 1)



if BO32<sub>0</sub> then CTR<sub>32:63</sub> ← CTR<sub>32:63</sub> - 1

ctr\_ok ← ¬BO32<sub>0</sub> | ((CTR<sub>32:63</sub> ≠ 0) ⊕ BO32<sub>1</sub>)

cond\_ok ← BO32<sub>0</sub> | (CR<sub>BI32+32</sub> ≡ BO32<sub>1</sub>)

if ctr\_ok & cond\_ok then

NIA ← (CIA + EXTS(BD15 || 0b0))<sub>32:63</sub>

else NIA ← CIA + 4

if LK=1 then LR ← CIA + 4

Let the BTEA be calculated as follows:

- For **e\_bc[l]**, let BTEA be the sum of the CIA and the sign-extended value of the BD15 instruction field concatenated with 0b0.



BO32 specifies any conditions that must be met for the branch to be taken, as defined in [Section 13.2.2: Branch instructions](#). The sum BI32+32 specifies the CR bit. Only CR[32–47] may be specified.

If the branch conditions are met, the BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA + 4 is placed into the LR.

Special Registers Altered: CTR (if BO32<sub>0</sub> = 1)  
LR (if LK = 1)

<b>se_bc</b>					BO16, BI16, BD8												
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		
1	1	1	0	0	BO16	BI16	BD8										

```

cond_ok ← (CRBI16+32 ≡ BO16)
if cond_ok then
    NIA ← (CIA + EXTS(BD8 || 0b0))32:63
else
    NIA ← CIA + 2
    
```

Let the BTEA be calculated as follows:

- For **se\_bc**, BTEA is the sum of the CIA and the sign-extended value of the BD8 instruction field concatenated with 0b0.

BO16 specifies any conditions that must be met for the branch to be taken, as defined in [Section 13.2.2: Branch instructions](#). The sum BI16+32 specifies CR bit; only CR[32–35] may be specified. If the branch conditions are met, the BTEA is the address of the next instruction to be executed.

Special Registers Altered: None

<b>_bclri</b>	VLE	User	<b>_bclri</b>
---------------	-----	------	---------------

Bit Clear Immediate

<b>se_bclri</b>					rX, UI5										
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	1	1	0	0	0	0	UI5					RX			

```

a ← UI5
b ← a1 || 0 || 31-a1
result32:63 ← GPR(RX) & b
GPR(RX) ← result32:63
    
```


For **se\_bclri**, the bit of GPR(rX) specified by the value of UI5 is cleared and all other bits in GPR(rX) remain unaffected.

Special Registers Altered: None

<b>_bctrx</b>	VLE	User	<b>_bctrx</b>
---------------	-----	------	---------------

Branch to Count Register [and Link]

**se\_bctr** (LK = 0)  
**se\_bctrl** (LK = 1)



$$NIA \leftarrow CTR_{32:62} \parallel 0b0$$

if LK=1 then LR ← CIA + 2

Let the BTEA be calculated as follows:

- For **se\_bctr**[I], let BTEA be bits 32–62 of the contents of the CTR concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA + 2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)


**\_bgeni**

VLE	User
-----	------

**\_bgeni**

Bit Generate Immediate

**se\_bgeni** rX,UI5



$$a \leftarrow UI5$$

$$b \leftarrow a0 \parallel 1 \parallel 31-a0$$

$$GPR(RX) \leftarrow b$$

For **se\_bgeni**, a constant value consisting of a single '1' bit surrounded by '0's is generated and the value is placed into GPR(rX). The position of the '1' bit is specified by the UI5 field.

Special Registers Altered: None


**\_blrx**

VLE	User
-----	------

**\_blrx**

Branch to Link Register [and Link]

**se\_blr** (LK = 0)  
**se\_blrl** (LK = 1)



$$NIA \leftarrow LR_{32:62} \parallel 0b0$$

if LK=1 then LR ← CIA + 2

Let the BTEA be calculated as follows:

- For **se\_blr**[I], let BTEA be bits 32–62 of the contents of the LR concatenated with 0b0.

The BTEA is the address of the next instruction to be executed.

If LK = 1, the sum CIA + 2 is placed into the LR.

Special Registers Altered: LR (if LK = 1)


**\_bmaski**

VLE	User
-----	------

**\_bmaski**

Bit Mask Generate Immediate

**se\_bmaski** rX,UI5



$a \leftarrow \text{UI5}$   
 if  $a = 0$  then  $b \leftarrow 32\text{1}$  else  $b \leftarrow 32\text{-}a0 \parallel a1$   
 $\text{GPR}(\text{RX}) \leftarrow b$

For **se\_bmaski**, a constant value consisting of a mask of low-order '1' bits that is zero-extended to 32 bits is generated, and the value is placed into GPR(rX). The number of low-order '1' bits is specified by the UI5 field. If UI5 is 0b00000, a value of all '1's is generated

Special Registers Altered: None


**\_bseti**

VLE	User
-----	------

**\_bseti**

Bit Set Immediate

**se\_bseti** rX,UI5



$a \leftarrow \text{UI5}$   
 $b \leftarrow a0 \parallel 1 \parallel 31\text{-}a0$   
 $\text{result}_{32:63} \leftarrow \text{GPR}(\text{RX}) \mid b$   
 $\text{GPR}(\text{RX}) \leftarrow \text{result}_{32:63}$

For **se\_bseti**, the bit of GPR(rX) specified by the value of UI5 is set, and all other bits in GPR(rX) remain unaffected.

Special Registers Altered: None


**\_btsti**

VLE	User
-----	------

**\_btsti**

Bit Test Immediate

**se\_btsti** rX,UI5




$a \leftarrow \text{UI5}$   
 $b \leftarrow a0 \parallel 1 \parallel 31\text{-}a0$   
 $c \leftarrow \text{GPR}(\text{RX}) \& b$   
 if  $c = 32\text{0}$  then  $d \leftarrow 0\text{b}001$  else  $d \leftarrow 0\text{b}010$

$$CR_{0:3} \leftarrow d \parallel XER_{SO}$$


For **se\_btsti**, the bit of GPR(**rX**) specified by the value of UI5 is tested for equality to '1'. The result of the test is recorded in the CR. EQ is set if the tested bit is clear, LT is cleared, and GT is set to the inverse value of EQ.

Special Registers Altered: CR[0–3]

<b>_cmp</b>	VLE	User	<b>_cmp</b>
Compare [Immediate]			
<b>e_cmp16i</b>	<b>rA,SI</b>		



**e\_cmpi** **crD32,rA,SCI8**





$a \leftarrow GPR(RA)_{32:63}$   
 if 'e\_cmpi' then  $b \leftarrow SCI8(F,SCL,UI8)$   
 if 'e\_cmp16i' then  $b \leftarrow EXTS(SI_{0:4} \parallel SI_{5:15})$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 if 'e\_cmpi' then  $CR_{4 \times CRD32+32:4 \times CRD32+35} \leftarrow c \parallel XER_{SO}$  // only CR0-CR3  
 if 'e\_cmp16i' then  $CR_{32:35} \leftarrow c \parallel XER_{SO}$  // only CR0

If **e\_cmpi**, GPR(**rA**) contents are compared with the value of SCI8, treating operands as signed integers.

If **e\_cmp16i**, GPR(**rA**) contents are compared with the sign-extended value of the SI field, treating operands as signed integers.

The result of the comparison is placed into CR field **crD** (**crD32**). For **e\_cmpi**, only CR0–CR3 may be specified. For **e\_cmp16i**, only CR0 may be specified.

Special Registers Altered: CR field **crD** (**crD32**) (CR0 for **e\_cmp16i**)



$a \leftarrow GPR(RX)_{32:63}$   
 if 'se\_cmpi' then  $b \leftarrow {}^{27}0 \parallel UI5$

if 'se\_cmp' then  $b \leftarrow \text{GPR}(\text{rY})_{32:63}$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $\text{CR}_{0:3} \leftarrow c \parallel \text{XER}_{\text{SO}}$

If **se\_cmp**, the contents of GPR(rX) are compared with the contents of GPR(rY), treating the operands as signed integers. The result of the comparison is placed into CR field 0.

If **se\_cmpi**, the contents of GPR(rX) are compared with the value of the zero-extended UI5 field, treating the operands as signed integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]

**\_cmph**

VLE	User
-----	------

**\_cmph**

Compare Halfword [Immediate]

**e\_cmph** **crD,rA,rB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRD	/			RA				RB			0	0	0	0	0	0	1	1	1	0	/				

$a \leftarrow \text{EXTS}(\text{GPR}(\text{RA})_{48:63})$   
 $b \leftarrow \text{EXTS}(\text{GPR}(\text{RB})_{48:63})$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $\text{CR}_{4 \times \text{CRD} + 32 : 4 \times \text{CRD} + 35} \leftarrow c \parallel \text{XER}_{\text{SO}}$

For **e\_cmph**, the contents of the low-order 16 bits of GPR(rA) and GPR(rB) are compared, treating the operands as signed integers. The result of the comparison is placed into CR field CRD.

Special Registers Altered: CR field CRD

**se\_cmph** **rX,rY**


0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	1	1	1	0			RY				RX	

$a \leftarrow \text{EXTS}(\text{GPR}(\text{RX})_{48:63})$   
 $b \leftarrow \text{EXTS}(\text{GPR}(\text{RY})_{48:63})$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $\text{CR}_{0:3} \leftarrow c \parallel \text{XER}_{\text{SO}}$

For **se\_cmph**, the contents of the low-order 16 bits of GPR(**rX**) and GPR(**rY**) are compared, treating the operands as signed integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]

**e\_cmph16i**                      **rA,SI**




```

a ← EXTS(GPR(RA)48:63)
b ← EXTS(SI0:4 || SI5:15)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR32:35 ← c || XERSO // only CR0
    
```

The contents of the lower 16-bits of GPR(**rA**) are sign-extended and compared with the sign-extended value of the SI field, treating the operands as signed integers.


The result of the comparison is placed into CR0.

Special Registers Altered: CR0



Compare Halfword Logical [Immediate]

**e\_cmphl**                      **crD,rA,rB**




```

a ← EXTZ(GPR(RA)48:63)
b ← EXTZ(GPR(RB)48:63)
if a < b then c ← 0b100
if a > b then c ← 0b010
if a = b then c ← 0b001
CR4×CRD+32:4×CRD+35 ← c || XERSO
    
```

For **e\_cmphl**, the contents of the low-order 16 bits of GPR(**rA**) and GPR(**rB**) are compared, treating the operands as unsigned integers. The result of the comparison is placed into CR field CRD.

Special Registers Altered: CR field CRD

**se\_cmphl**                      **rX,rY**



```

a ← GPR(RX)48:63
    
```

$b \leftarrow \text{GPR}(\text{RY})_{48:63}$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $\text{CR}_{0:3} \leftarrow c \parallel \text{XER}_{\text{SO}}$

For **se\_cmphi**, the contents of the low-order 16 bits of GPR(rX) and GPR(rY) are compared, treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]

**e\_cmphi16i**                      rA,UI

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	0	UI <sub>0:4</sub>					RA				1	0	1	1	1	UI <sub>5:15</sub>											

$a \leftarrow {}^{16}0 \parallel \text{GPR}(\text{RA})_{48:63}$   
 $b \leftarrow {}^{16}0 \parallel \text{UI}_{0:4} \parallel \text{UI}_{5:15}$   
 if  $a < b$  then  $c \leftarrow 0b100$   
 if  $a > b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $\text{CR}_{32:35} \leftarrow c \parallel \text{XER}_{\text{SO}}$  // only CR0

The contents of the lower 16-bits of GPR(rA) are zero-extended and compared with the zero-extended value of the UI field, treating the operands as unsigned integers.

The result of the comparison is placed into CR0.

Special Registers Altered: CR0

\_\_cmpli                      VLE    User                      \_\_cmpli  
 \_\_cmpli \_\_cmpli \_\_cmpli \_\_cmpli

Compare Logical [Immediate]

**e\_cmpli16i**                      rA,UI

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	0	UI <sub>0:4</sub>					RA				1	0	1	0	1	UI <sub>5:15</sub>											

**e\_cmpli**                                      crD32,rA,SCI8

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	1	0	0	0	1	CRD32			RA				1	0	1	0	1	F	SCL	UI8								

$a \leftarrow \text{GPR}(\text{RA})_{32:63}$   
 if 'e\_cmpli' then  $b \leftarrow \text{SCI8}(\text{F}, \text{SCL}, \text{UI8})$   
 if 'e\_cmpli16i' then  $b \leftarrow {}^{16}0 \parallel \text{UI}_{0:4} \parallel \text{UI}_{5:15}$   
 if  $a <_u b$  then  $c \leftarrow 0b100$   
 if  $a >_u b$  then  $c \leftarrow 0b010$



if  $a = b$  then  $c \leftarrow 0b001$   
 if 'e\_cmpli' then  $CR_{4 \times CRD32+32:4 \times CRD32+35} \leftarrow c \parallel XER_{SO}$  // only CR0-CR3  
 if 'e\_cmp16i' then  $CR_{32:35} \leftarrow c \parallel XER_{SO}$  // only CR0

If **e\_cmpli**, the contents of bits 32–63 of GPR(rA) are compared with the value of SC18, treating the operands as unsigned integers.

L must be 0 for 32-bit implementations

If **e\_cmpli16i**, the contents of GPR(rA) are compared with the zero-extended value of the UI field, treating the operands as unsigned integers.

The result of the comparison is placed into CR field CRD (CRD32). For **e\_cmpli**, only CR0–CR3 may be specified. For **e\_cmpli16i**, only CR0 may be specified.

Special Registers Altered: CR field CRD (CRD32) (CR0 for **e\_cmpli16i**)

<b>se_cmpl</b>															<b>rX,rY</b>		
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		
0	0	0	0	1	1	0	1	RY		RX							

<b>se_cmpli</b>															<b>rX,OIMM</b>		
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15		
0	0	1	0	0	0	1	OIM5 <sup>(1)</sup>		RX								

1. OIMM = OIM5 +1

$a \leftarrow GPR(RX)_{32:63}$   
 if 'se\_cmpli' then  $b \leftarrow {}^{27}0 \parallel OFFSET(OIM5)$   
 if 'se\_cmpl' then  $b \leftarrow GPR(RY)_{32:63}$   
 if  $a <_u b$  then  $c \leftarrow 0b100$   
 if  $a >_u b$  then  $c \leftarrow 0b010$   
 if  $a = b$  then  $c \leftarrow 0b001$   
 $CR_{0:3} \leftarrow c \parallel XER_{SO}$

If **se\_cmpl**, the contents of GPR(rX) are compared with the contents of GPR(rY), treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.

If **se\_cmpli**, the contents of GPR(rX) are compared with the value of the zero-extended offset value of the OIM5 field (a final value in the range 1–32), treating the operands as unsigned integers. The result of the comparison is placed into CR field 0.

Special Registers Altered: CR[0–3]

<b>__crand</b>															VLE		User		<b>__crand</b>														
<b>__crand __crand __crand __crand</b>																																	

Condition Register AND

<b>e_crاند</b>															<b>crbD,crbA,crbB</b>																
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	CRBD		CRBA				CRBB				0	1	0	0	0	0	0	0	0	0	0	0	0	1	/	



$$CR_{BT+32} \leftarrow CR_{BA+32} \& CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ANDed with the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

Condition Register AND with Complement

**e\_crandc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1																	0	0	1	0	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \& \neg CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ANDed with the one's complement of the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

CR Equivalent

**e\_creqv** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1																	0	1	0	0	1	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \oplus CR_{BB+32}$$

The content of bit CRBA+32 of the CR is XORed with the content of bit CRBB+32 of the CR, and the one's complement of result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_crnand** VLE User **\_crnand**  
**\_crnand** **\_crnand** **\_crnand** **\_crnand** **\_crnand** **\_crnand**

Condition Register NAND

**e\_crnand** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1																	0	0	1	1	1	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow \neg(CR_{BA+32} \& CR_{BB+32})$$

The content of bit CRBA+32 of the CR is ANDed with the content of bit CRBB+32 of the CR, and the one's complement of the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_crnor** VLE User **\_crnor**

Condition Register NOR

**e\_crnor** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
0	1	1	1	1	1																	0	0	0	0	1	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow \neg(CR_{BA+32} | CR_{BB+32})$$

The content of bit CRBA+32 of the CR is ORed with the content of bit CRBB+32 of the CR, and the one's complement of the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_\_cror** VLE User **\_\_cror**

Condition Register OR

**e\_cror** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	1	1	1	1	1	CRBD						CRBA						CRBB						0	1	1	1	0	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \mid CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ORed with the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_\_cror** VLE User **\_\_cror**  
**\_\_cror \_\_cror \_\_cror \_\_cror** \_

Condition Register OR with Complement

**e\_crorc** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	1	1	1	1	1	CRBD						CRBA						CRBB						0	1	1	0	1	0	0	0	0	0	1	/

$$CR_{BT+32} \leftarrow CR_{BA+32} \mid \neg CR_{BB+32}$$

The content of bit CRBA+32 of the CR is ORed with the one's complement of the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_\_crxor** VLE User **\_\_crxor**  
**\_\_crxor \_\_crxor \_\_crxor \_\_crxor** \_

Condition Register XOR

**e\_crxor** **crbD,crbA,crbB**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31				
0	1	1	1	1	1	CRBD						CRBA						CRBB						0	0	1	1	0	0	0	0	0	0	1	/

$$CR_{crbD+32} \leftarrow CR_{BA+32} \oplus CR_{BB+32}$$

The content of bit CRBA+32 of the CR is XORed with the content of bit CRBB+32 of the CR, and the result is placed into bit CRBD+32 of the CR.

Special Registers Altered: CR

**\_\_extsbx** VLE User **\_\_extsbx**

Extend Sign (Byte | Halfword)

<b>se_extsb</b>						<b>rX</b>									
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	1	1	0	1	RX			

<b>se_extsh0.</b>						<b>rX</b>									
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	1	1	1	1	RX			

if se\_extsb then n ← 56

if se\_extsh then n ← 48

if 'extsw' then n ← 32

if Rc=1 then do

LT ← GPR(RS)<sub>n:63</sub> < 0

GT ← GPR(RS)<sub>n:63</sub> > 0

EQ ← GPR(RS)<sub>n:63</sub> = 0

CR0 ← LT || GT || EQ || SO

s ← GPR(RS or RX)<sub>n</sub>

GPR(RA or RX) ← <sup>n-32</sup>s || GPR(RS or RX)<sub>n:63</sub>

For **se\_extsb**, the contents of bits 56–63 of GPR(rX) are placed into bits 56–63 of GPR(rX). Bit 56 of the contents of GPR(rX) is copied into bits 32–55 of GPR(rX).

For **se\_extsh**, the contents of bits 48–63 of GPR(rX) are placed into bits 48–63 of GPR(rX). Bit 48 of the contents of GPR(rX) is copied into bits 32–47 of GPR(rX).

Special Registers Altered: CR0 (if Rc=1)

<b>_extzx</b>	VLE	User	<b>extzx</b>
---------------	-----	------	--------------

Extend Zero (Byte | Halfword)

<b>se_extzb</b>						<b>rX</b>									
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	1	1	0	0	RX			

<b>se_extzh</b>						<b>rX</b>									
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	1	1	1	0	RX			

if 'se\_extzb' then n ← 56

if 'se\_extzh' then n ← 48

GPR(RX) ← <sup>n-32</sup>0 || GPR(RX)<sub>n:63</sub>

For **se\_extzb**, the contents of bits 56–63 of GPR(rX) are placed into bits 56–63 of GPR(rX). Bits 32–55 of GPR(rX) are cleared.

For **se\_extzh**, the contents of bits 48–63 of GPR(rX) are placed into bits 48–63 of GPR(rX). Bits 32–47 of GPR(rX) are cleared.

Special Registers Altered: None

**\_illegal**

VLE	User
-----	------

**\_illegal**

Illegal

**se\_illegal**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

SRR1 ← MSR

SRR0 ← CIA

NIA ← IVPR<sub>32:47</sub> || IVOR6<sub>48:59</sub> || 0b0000

MSR<sub>WE,EE,PR,IS,DS,FP,FE0,FE1</sub> ← 0b0000\_0000

**se\_illegal** is used to request an illegal instruction exception. A program interrupt is generated. The contents of the MSR are copied into SRR1 and the address of the **se\_illegal** instruction is placed into SRR0.

MSR[WE,EE,PR,IS,DS,FP,FE0,FE1] are cleared.

The interrupt causes the next instruction to be fetched from address IVPR[32–47]||IVOR6[48–59]||0b0000

This instruction is context synchronizing.

Special Registers Altered: SRR0 SRR1 MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]

**\_isync**

VLE	User
-----	------

**\_isync**

Instruction Synchronize

**se\_isync**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

The **se\_isync** instruction provides an ordering function for the effects of all instructions executed by the processor executing the **se\_isync** instruction. Executing an **se\_isync** instruction ensures that all instructions preceding the **se\_isync** instruction have completed before the **se\_isync** instruction completes, and that no subsequent instructions are initiated until after the **se\_isync** instruction completes. It also causes any prefetched instructions to be discarded, with the effect that subsequent instructions are fetched and executed in the context established by the instructions preceding the **se\_isync** instruction.

The **se\_isync** instruction may complete before memory accesses associated with instructions preceding the **se\_isync** instruction have been performed.

This instruction is context synchronizing (see Book E). It has identical semantics to Book E **isync**, just a different encoding.




Special Registers Altered: None

**\_lbzx**

VLE	User
-----	------

**\_lbzx**

Load Byte and Zero [with Update] [Indexed]



if (RA=0 & !se\_lbz) then a ← 32'0 else a ← GPR(RA or RX)

if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>

if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>

if SD4-mode then EA ← (a + (28'0 || SD4))<sub>32:63</sub>

GPR(RD or RZ) ← 24'0 || MEM(EA,1)

if e\_lbzu then GPR(RA) ← EA

Let the EA be calculated as follows:

- For **e\_lbz** and **e\_lbzu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_lbz**, let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field.

The byte in memory addressed by EA is loaded into bits 56–63 of GPR(rD or rZ). Bits 32–55 of GPR(rD or rZ) are cleared.



If **e\_lbzu**, EA is placed into GPR(rA).

If **e\_lbzu** and rA = 0 or rA = rD, the instruction form is invalid.

Special Registers Altered: None



Load Halfword Algebraic [with Update] [Indexed]



if RA=0 then a ← 32'0 else a ← GPR(RA)

if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>

if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>

$GPR(RD) \leftarrow EXTS(MEM(EA,2))_{32:63}$

if **e\_lhau** then  $GPR(RA) \leftarrow EA$

Let the EA be calculated as follows:

- For **e\_lha** and **e\_lhau**, let EA be the sum of the contents of  $GPR(rA)$ , or 32 0s if  $rA = 0$ , and the sign-extended value of the D or D8 instruction field.

The half word in memory addressed by EA is loaded into bits 48–63 of  $GPR(rD)$ . Bits 32–47 of  $GPR(rD)$  are filled with a copy of bit 0 of the loaded half word.

If **e\_lhau**, EA is placed into  $GPR(rA)$ .

If **e\_lhau** and  $rA = 0$  or  $rA = rD$ , the instruction form is invalid.

Special Registers Altered: None


**\_lhzx**

VLE	User
-----	------


**\_lhzx**

Load Halfword and Zero [with Update] [Indexed]


**e\_lhz**  $rD, D(rA)$  (D-mode)



**se\_lhz**  $rZ, SD4(rX)$  (SD4-mode)



**e\_lhzu**  $rD, D8(rA)$  (D8-mode)



if  $(RA=0 \ \& \ !se\_lhz)$  then  $a \leftarrow {}^{32}0$  else  $a \leftarrow GPR(RA \text{ or } RX)$

if D-mode then  $EA \leftarrow (a + EXTS(D))_{32:63}$

if D8-mode then  $EA \leftarrow (a + EXTS(D8))_{32:63}$

if SD4-mode then  $EA \leftarrow (a + ({}^{27}0 \ || \ SD4 \ || \ 0))_{32:63}$

$GPR(RD \text{ or } RZ) \leftarrow {}^{16}0 \ || \ MEM(EA,2)$

if **e\_lhzu** then  $GPR(RA) \leftarrow EA$

Let the EA be calculated as follows:

- For **e\_lhz** and **e\_lhzu**, let EA be the sum of the contents of  $GPR(rA)$ , or 32 0s if  $rA = 0$ , and the sign-extended value of the D or D8 instruction field.
- For **se\_lhz** let EA be the sum of the contents of  $GPR(rX)$  and the zero-extended value of the SD4 instruction field shifted left by 1 bit.

The half word in memory addressed by EA is loaded into bits 48–63 of  $GPR(rD)$ . Bits 32–47 of  $GPR(rD)$  are cleared.

If **e\_lhzu**, EA is placed into  $GPR(rA)$ .

If **e\_lhzu** and  $rA = 0$  or  $rA = rD$ , the instruction form is invalid.

Special Registers Altered: None


\_lix

VLE	User
-----	------

\_lix

Load Immediate [Shifted]

**e\_li** rD,LI20 (LI20-mode)




$$LI20 \leftarrow LI20_{0:3} \parallel LI20_{4:8} \parallel LI20_{9:19}$$

$$GPR(RD) \leftarrow EXTS(LI20)$$

For **e\_li**, the sign-extended LI20 field is placed into GPR(rD).

Special Registers Altered: None

**e\_lis** rD,UI




$$UI \leftarrow UI_{0:4} \parallel UI_{5:15}$$

$$GPR(RD) \leftarrow UI \parallel 16_0$$

For **e\_lis**, the UI field is concatenated on the right with 16 0's and placed into GPR(rD).

Special Registers Altered: None

**se\_li** rX,UI7



$$GPR(RX) \leftarrow 25_0 \parallel UI7$$

For **se\_li**, the zero-extended UI7 field is placed into GPR(rX).

Special Registers Altered: None


\_lmw

VLE	User
-----	------

\_lmw

Load Multiple Word

**e\_lmw** rD,D8(rA)



$$\text{if } RA=0 \text{ then } EA \leftarrow EXTS(D8)_{32:63}$$

$$\text{else } EA \leftarrow (GPR(RA)+EXTS(D8))_{32:63}$$

$$r \leftarrow RD$$

$$\text{do while } r \leq 31$$

$$\quad GPR(r) \leftarrow MEM(EA,4)$$

$$\quad r \leftarrow r + 1$$

$$\quad EA \leftarrow (EA+4)_{32:63}$$

Let the EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D8 instruction field.



Let  $n = (32-rD)$ .  $n$  consecutive words starting at EA are loaded into bits 32–63 of registers GPR( $rD$ ) through GPR(31).

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined. If  $rA$  is in the range of registers to be loaded, including the case in which  $rA = 0$ , the instruction form is invalid.

Special Registers Altered: None


**\_lwz**

VLE	User
-----	------


**\_lwz**

Load Word and Zero [with Update] [Indexed]


**e\_lwz** **rD,D(rA)** (D-mode)



**se\_lwz** **rZ,SD4(rX)** (SD4-mode)



**e\_lwzu** **rD,D8(rA)** (D8-mode)



if ( $RA=0$  & ! $se\_lwz$ ) then  $a \leftarrow 32'0$  else  $a \leftarrow GPR(RA$  or  $RX)$

if D-mode then  $EA \leftarrow (a + EXTS(D))_{32:63}$

if D8-mode then  $EA \leftarrow (a + EXTS(D8))_{32:63}$

if SD4-mode then  $EA \leftarrow (a + (2^6'0 \parallel SD4 \parallel 2^0))_{32:63}$

$GPR(RD$  or  $RZ) \leftarrow MEM(EA,4)$

if  $e\_lwzu$  then  $GPR(RA) \leftarrow EA$

Let the EA be calculated as follows:

- For **e\_lwz** and **e\_lwzu**, let EA be the sum of the contents of GPR( $rA$ ), or 32 0s if  $rA = 0$ , and the sign-extended value of the D or D8 instruction field.
- For **se\_lwz** let EA be the sum of the contents of GPR( $rX$ ) and the zero-extended value of the SD4 instruction field shifted left by 2 bits.

The word in memory addressed by the EA is loaded into bits 32–63 of GPR( $rD$ ).

If **e\_lwzu**, EA is placed into GPR( $rA$ ).

If **e\_lwzu** and  $rA = 0$  or  $rA = rD$ , the instruction form is invalid.

Special Registers Altered: None


**\_mcrf**

VLE	User
-----	------

**\_mcrf**

Move CR Field

**e\_mcrf** **crD,crS**



$$CR_{4 \times CRD + 32 : 4 \times CRD + 35} \leftarrow CR_{4 \times CRS + 32 : 4 \times CRS + 35}$$

The contents of field **crS** (bits  $4 \times CRS + 32$  through  $4 \times CRS + 35$ ) of the CR are copied to field **crD** (bits  $4 \times CRD + 32$  through  $4 \times CRD + 35$ ) of the CR.

Special Registers Altered: CR

**\_mfar**

VLE	User
-----	------

**\_mfar**

Move from Alternate Register

<b>se_mfar</b>								<b>rX,arY</b>							
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	1	1	ARY	RX						

$$GPR(RX) \leftarrow GPR(ARY)$$

For **se\_mfar**, the contents of GPR(**arY**) are placed into GPR(**rX**). **arY** specifies a GPR in the range R8–R23. The encoding 0000 specifies R8, 0001 specifies R9, ..., 1111 specifies R23.

Special Registers Altered: None

**\_mfctr**

VLE	User
-----	------

**\_mfctr**

Move From Count Register

<b>se_mfctr</b>								<b>rX</b>								
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	0	0	0	0	0	0	0	0	1	0	1	0	RX		

$$GPR(RX) \leftarrow CTR$$

The CTR contents are placed into bits 32–63 of GPR(**rX**).

Special Registers Altered: None

**\_mflr**

VLE	User
-----	------

**\_mflr**

Move From Link Register

<b>se_mflr</b>								<b>rX</b>								
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	0	0	0	0	0	0	0	0	1	0	0	0	RX		

$$GPR(RX) \leftarrow LR$$

The LR contents are placed into bits 32–63 of GPR(**rX**).

Special Registers Altered: None

**\_mr**

VLE	User
-----	------

**\_mr**

Move Register

<b>se_mr</b>								<b>rX,rY</b>							
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	1	RY	RX						

$$GPR(RX) \leftarrow GPR(RY)$$

For **se\_mr**, the contents of GPR(**rY**) are placed into GPR(**rX**).


Special Registers Altered: None

**\_mtar**

VLE	User
-----	------

**\_mtar**

Move to Alternate Register



$$GPR(ARX) \leftarrow GPR(RY)$$

For **se\_mtar**, the contents of GPR(**rY**) are placed into GPR(**arX**). **arX** specifies a GPR in the range R8–R23. The encoding 0000 specifies R8, 0001 specifies R9, ..., 1111 specifies R23.


Special Registers Altered: None

**\_mtctr**

VLE	User
-----	------

**\_mtctr**

Move To Count Register



$$CTR \leftarrow GPR(RX)$$

The contents of bits 32–63 of GPR(**rX**) are placed into the CTR.


Special Registers Altered: CTR

**\_mtlr**

VLE	User
-----	------

**\_mtlr**

Move To Link Register



$$LR \leftarrow GPR(RX)$$

The contents of bits 32–63 of GPR(**rX**) are placed into the LR.

Special Registers Altered: LR


**\_mullix**

VLE	User
-----	------

**\_mullix**

Multiply Low [2 operand] Immediate

**e\_mulli** **rD,rA,SCI8**



$$imm \leftarrow SCI8(F,SCL,UI8)$$

$$prod_{0:63} \leftarrow GPR(RA) \times imm$$


$$GPR(RD) \leftarrow prod_{32:63}$$

Bits 32–63 of the 64-bit product of the contents of GPR(**rA**) and the value of SCI8 are placed into GPR(**rD**).

Both operands and the product are interpreted as signed integers.

Special Registers Altered: None

**e\_mull2i** **rA,SI**



$$\text{prod}_{0:63} \leftarrow \text{GPR}(\text{RA}) \times \text{EXTS}(\text{SI}_{0:4} \parallel \text{SI}_{5:15})$$

$$\text{GPR}(\text{RA}) \leftarrow \text{prod}_{32:63}$$

Bits 32–63 of the 64-bit product of the contents of GPR(**rA**) and the sign-extended value of the SI field are placed into GPR(**rA**).

Both operands and the product are interpreted as signed integers.

Special Registers Altered: None


**\_mullwx**

VLE	User
-----	------

**\_mullwx**

Multiply Low Word

**se\_mullw** **rX,rY**



$$\text{prod}_{0:63} \leftarrow \text{GPR}(\text{RX})_{32:63} \times \text{GPR}(\text{RY})_{32:63}$$

$$\text{GPR}(\text{RX}) \leftarrow \text{prod}_{32:63}$$

Bits 32–63 of the 64-bit product of the contents of bits 32–63 of GPR(**rX**) and the contents of bits 32–63 of GPR(**rY**) is placed into GPR(**rX**).

Special Registers Altered: None


**\_negx**

VLE	User
-----	------

**\_negx**

Negate

**se\_neg** **rX**



$$\text{result}_{32:63} \leftarrow \neg \text{GPR}(\text{RX}) + 1$$

$$\text{GPR}(\text{RX}) \leftarrow \text{result}_{32:63}$$

The sum of the one's complement of the contents of GPR(**rX**) and 1 is placed into GPR(**rX**).

If bits 32–63 of GPR(**rX**) contain the most negative 32-bit number (0x8000\_0000), bits 32–63 of the result contain the most negative 32-bit number


Special Registers Altered: None

**\_notx**

VLE	User
-----	------

**\_notx**

NOT








result<sub>32:63</sub> ← ¬GPR(RX)

GPR(RX) ← result<sub>32:63</sub>

The contents of GPR(rX) are inverted.

Special Registers Altered: None



if 'e\_ori[.]' then b ← SCI8(F,SCL,UI8)

if 'e\_or2i' then b ← <sup>16</sup>0 || UI<sub>0:4</sub> || UI<sub>5:15</sub>

if 'e\_or2is' then b ← UI<sub>0:4</sub> || UI<sub>5:15</sub> || <sup>16</sup>0

if 'se\_or' then b ← GPR(RB)

result<sub>0:63</sub> ← GPR(RS or RD or RX) | b

if Rc=1 then do

LT ← result<sub>32:63</sub> < 0

GT ← result<sub>32:63</sub> > 0

EQ ← result<sub>32:63</sub> = 0

CR0 ← LT || GT || EQ || SO

GPR(RA or RD or RX) ← result

For **e\_ori[.]**, the contents of GPR(rS) are ORed with the value of SCI8.

For **e\_or2i**, the contents of GPR(rD) are ORed with <sup>16</sup>0 || UI.

For **e\_or2is**, the contents of GPR(rD) are ORed with UI || <sup>16</sup>0.

For **se\_or**, the contents of GPR(rX) are ORed with the contents of GPR(rY).

The result is placed into GPR(rA or rX).

The preferred 'no-op' (an instruction that does nothing) is:

```
e_ori 0,0,0
```

Special Registers Altered: CR0 (if Rc = 1)

**\_rfci** VLE Supervisor **\_rfci**

Return From Critical Interrupt

**se\_rfci**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1

MSR ← CSRR1

NIA ← CSRR0<sub>0:62</sub> || 0b0

The **se\_rfci** instruction is used to return from a critical class interrupt, or as a means of establishing a new context and synchronizing on that new context simultaneously.

The contents of CSRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched, under control of the new MSR value, from the address CSRR0[32–62]||0b0. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism (see Book E) is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in CSRR0 at the time of the execution of the **se\_rfci**).

Execution of this instruction is privileged and restricted to supervisor mode.

Execution of this instruction is context synchronizing.

Special Registers Altered: MSR

**\_rfi** VLE Supervisor **\_rfi**

Return From Interrupt

**se\_rfi**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0

MSR ← SRR1

NIA ← SRR0<sub>0:62</sub> || 0b0

The **se\_rfi** instruction is used to return from a non-critical class interrupt, or as a means of simultaneously establishing a new context and synchronizing on that new context.

The contents of SRR1 are placed into the MSR. If the new MSR value does not enable any pending exceptions, then the next instruction is fetched under control of the new MSR value from the address SRR0[32–62]||0b0. If the new MSR value enables one or more pending exceptions, the interrupt associated with the highest priority pending exception is generated; in this case the value placed into SRR0 or CSRR0 by the interrupt processing mechanism (see Book E) is the address of the instruction that would have been executed next had the interrupt not occurred (that is, the address in SRR0 at the time of the execution of the **se\_rfi**).

Execution of this instruction is privileged and restricted to supervisor mode.

Execution of this instruction is context synchronizing.

Special Registers Altered: MSR

	<b>_rlw</b>		VLE	User		<b>_rlw</b>																									
Rotate Left Word [Immediate]																															
		<b>e_rlwi</b>				<b>rA,rS,rB</b>	(Rc = 0)																								
		<b>e_rlwi.</b>				<b>rA,rS,rB</b>	(Rc = 1)																								
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	RS		RA			RB			0	1	0	0	0	1	1	0	0	0	0	0	0	0	0	0	Rc	
Rotate Left Word Immediate then Mask Insert																															
		<b>e_rlwi</b>				<b>rA,rS,SH</b>	(Rc = 0)																								
		<b>e_rlwi.</b>				<b>rA,rS,SH</b>	(Rc = 1)																								
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	RS		RA			SH			0	1	0	0	1	1	1	0	0	0	0	0	0	0	0	0	Rc	

```

if 'e_rlwi[.]' then n ← GPR(RB)59:63
else n ← SH
result32:63 ← ROTL32(GPR(RS)32:63,n)
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA) ← result32:63
    
```

If **e\_rlwi[.]**, let the shift count *n* be the contents of bits 59–63 of GPR(**rB**).

If **e\_rlwi[.]**, let the shift count *n* be SH.

The contents of GPR(**rS**) are rotated<sub>32</sub> left *n* bits. The rotated data is placed into GPR(**rA**).

Special Registers Altered: CR0 (if Rc = 1)

	<b>_rlwimi</b>		VLE	User		<b>_rlwimi</b>																									
Rotate Left Word Immediate then Mask Insert																															
		<b>e_rlwimi</b>				<b>rA,rS,SH,MB,ME</b>																									
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	1	RS		RA			SH			MB			ME			0											

```

n ← SH
b ← MB+32
e ← ME+32
r ← ROTL32(GPR(RS)32:63,n)
m ← MASK(b,e)
result32:63 ← r&m | GPR(RA)&~m
GPR(RA) ← result32:63
    
```





Let the shift count  $n$  be the value SH.

The contents of GPR( $rS$ ) are rotated<sub>32</sub> left  $n$  bits. A mask is generated having 1 bits from bit MB+32 through bit ME+32 and 0 bits elsewhere. The rotated data are inserted into GPR( $rA$ ) under control of the generated mask (if a mask bit is 1 the associated bit of the rotated data is placed into the target register, and if the mask bit is 0 the associated bit in the target register remains unchanged).

Special Registers Altered: None

**\_rlwinm**

VLE	User
-----	------

**\_rlwinm**

Rotate Left Word Immediate then AND with Mask

**e\_rlwinm**  $rA,rS,SH,MB,ME$

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	0	1	RS					RA					SH					MB					ME					1

- $n \leftarrow SH$
- $b \leftarrow MB+32$
- $e \leftarrow ME+32$
- $r \leftarrow ROTL_{32}(GPR(RS)_{32:63},n)$
- $m \leftarrow MASK(b,e)$
- $result_{32:63} \leftarrow r \& m$
- $GPR(RA) \leftarrow result_{32:63}$

Let the shift count  $n$  be SH.

The contents of GPR( $rS$ ) are rotated<sub>32</sub> left  $n$  bits. A mask is generated having 1 bits from bit MB+32 through bit ME+32 and 0 bits elsewhere. The rotated data are ANDed with the generated mask and the result is placed into GPR( $rA$ ).

Special Registers Altered: None

**\_sc**

VLE	User
-----	------

**\_sc**

System Call

**se\_sc**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0

- $SRR1 \leftarrow MSR$
- $SRR0 \leftarrow CIA+2$
- $NIA \leftarrow IVPR_{32:47} \parallel IVOR_{48:59} \parallel 0b0000$
- $MSR_{WE,EE,PR,IS,DS,FP,FE0,FE1} \leftarrow 0b0000\_0000$

**se\_sc** is used to request a system service. A system call interrupt is generated. The contents of the MSR are copied into SRR1 and the address of the instruction after the **se\_sc** instruction is placed into SRR0.

$MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]$  are cleared.



The interrupt causes the next instruction to be fetched from the address

IVPR[32–47]||IVOR8[48–59]||0b0000

This instruction is context synchronizing.

Special Registers Altered: SRR0 SRR1 MSR[WE,EE,PR,IS,DS,FP,FE0,FE1]

**\_slwx**

VLE	User
-----	------

**\_slwx**

Shift Left Word [Immediate] [and Record]

**e\_slwi** **rA,rS,SH** (Rc = 0)

**e\_slwi.** **rA,rS,SH** (Rc = 1)

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	1	1	1	1	1	RS					RA					SH			0	0	0	0	1	1	1	0	0	0	Rc		

**se\_slw** **rX,rY**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	1	0	0	0	0	1	0	RY				RX			

**se\_slwi** **rX,UI5**

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	1	1	0	1	1	0	UI5				RX				

```

if 'e_slwi[.]' then n ← SH
if se_slw then n ← GPR(RY)58:63
if se_slwi then n ← UI5
r ← ROTL32(GPR(RS or RX)32:63,n)
if n<32 then m ← MASK(32,63-n)
else m ← 320
result32:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
    
```

Let the shift count *n* be the value specified by the contents of bits 58–63 of GPR(*rB* or *rY*), or by the value of the SH or UI5 field.

The contents of bits 32–63 of GPR(*rS* or *rX*) are shifted left *n* bits. Bits shifted out of position 32 are lost. Zeros are supplied to the vacated positions on the right. The 32-bit result is placed into bits 32–63 of GPR(*rA* or *rX*).

Shift amounts from 32 to 63 give a zero result.



Special Registers Altered: CR0 (if Rc = 1)

**\_srawx**

VLE	User
-----	------

**\_srawx**

Shift Right Algebraic Word [Immediate] [and Record]



```

if 'se_sraw' then n ← GPR(RY)59:63
if 'se_srawi' then n ← UI5
r ← ROTL32(GPR(RS or RX)32:63,32-n)
if ((se_sraw & GPR(RY)58=1) then m ← 320
else m ← MASK(n+32,63)
s ← GPR(RS or RX)32
result0:63 ← r&m | (32s)&¬m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
CA ← s & ((r&¬m)32:63≠0)
    
```

If **se\_sraw**, let the shift count *n* be the contents of bits 58–63 of GPR(**rY**).



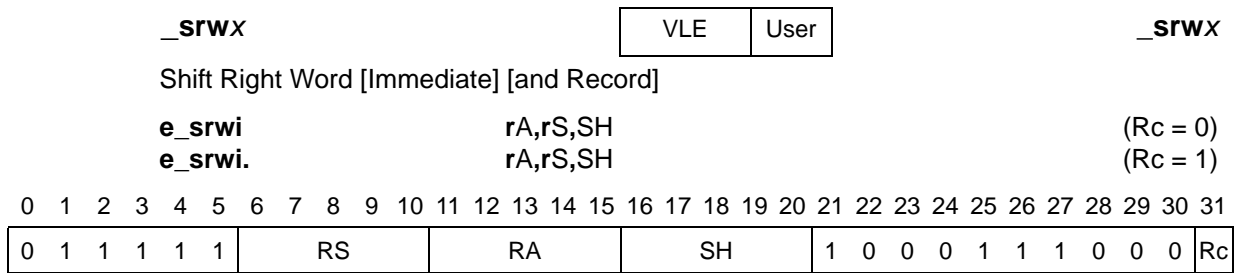
If **se\_srawi**, let the shift count *n* be the value of the UI5 field.

The contents of bits 32–63 of GPR(**rS** or **rX**) are shifted right *n* bits. Bits shifted out of position 63 are lost. Bit 32 of **rS** or **rX** is replicated to fill vacated positions on the left. The 32-bit result is placed into bits 32–63 of GPR(**rA** or **rX**).

CA is set if bits 32–63 of GPR(**rS** or **rX**) contain a negative value and any 1 bits are shifted out of bit position 63; otherwise CA is cleared.

A shift amount of zero causes GPR(**rA** or **rX**) to receive EXTS(GPR(**rS** or **rX**)<sub>32:63</sub>), and CA to be cleared. For **se\_sraw**, shift amounts from 32 to 63 give a result of 64 sign bits, and cause CA to receive bit 32 of the contents of GPR(**rS** or **rX**) (that is, sign bit of GPR(**rS** or **rX**)<sub>32:63</sub>).

Special Registers Altered: CA  
CR0 (if Rc = 1)



```

n ← GPR(RB)59:63

if 'e_srwi[.]' then n ← SH
if 'se_srw' then n ← GPR(RY)59:63
if 'se_srwi' then n ← UI5
r ← ROTL32(GPR(RS or RX)32:63,32-n)
if ((se_srw & GPR(RY)58=1) then m ← 320
else m ← MASK(n+32,63)
result32:63 ← r & m
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA or RX) ← result32:63
    
```

If **e\_srwi**, let the shift count *n* be the value of the SH field.

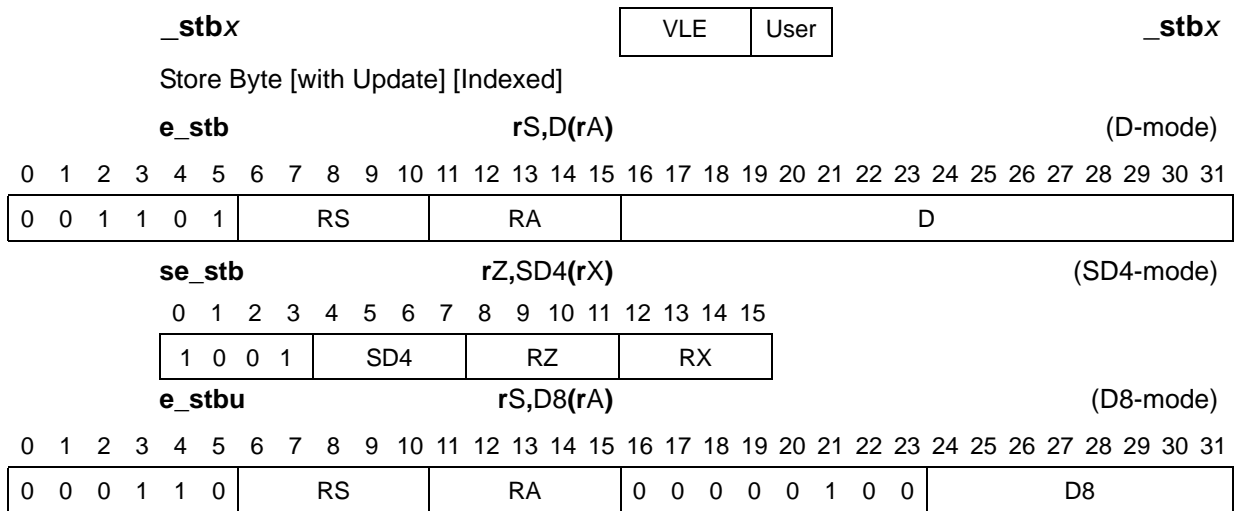
If **se\_srw**, let the shift count *n* be the contents of bits 58–63 of GPR(rY).

If **se\_srwi**, let the shift count *n* be the value of the UI5 field.

The contents of bits 32–63 of GPR(rS or rX) are shifted right *n* bits. Bits shifted out of position 63 are lost. Zeros are supplied to the vacated positions on the left. The 32-bit result is placed into bits 32–63 of GPR(rA or rX).

Shift amounts from 32 to 63 give a zero result.

Special Registers Altered: CR0 (if Rc = 1)



if (RA=0 & lse\_stb) then a ← <sup>32</sup>0 else a ← GPR(RA or RX)

if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>

if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>

if SD4-mode then EA ← (a + (<sup>28</sup>0 || SD4))<sub>32:63</sub>

MEM(EA,1) ← GPR(RS or RZ)<sub>56:63</sub>

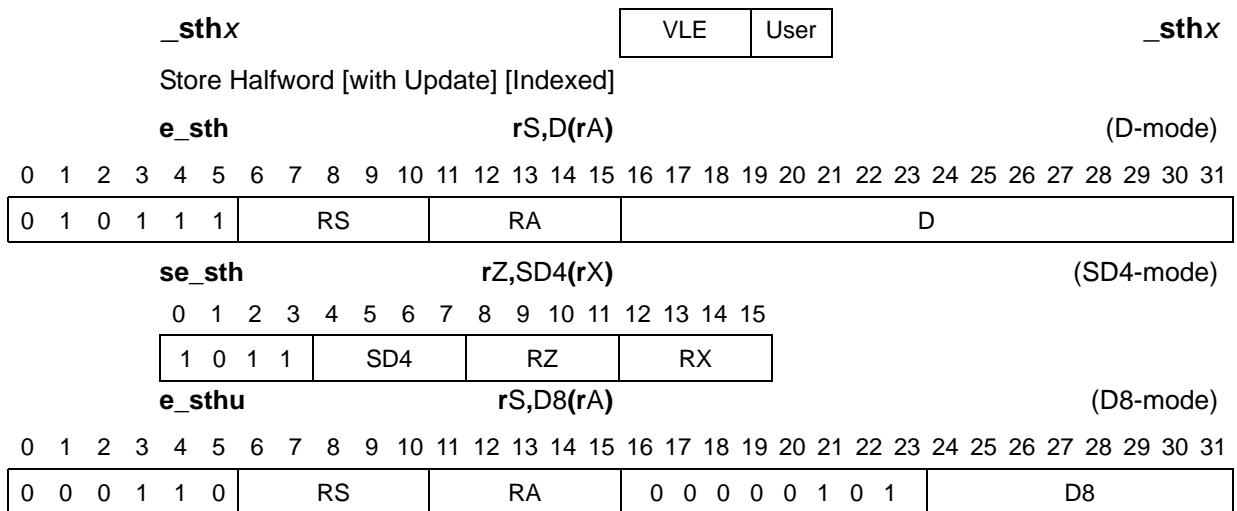
if e\_stbu then GPR(RA) ← EA

Let the EA be calculated as follows:

- For **e\_stb** and **e\_stbu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_stb**, let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field.

The contents of bits 56–63 of GPR(rS) are stored into the byte in memory addressed by EA.

- If **e\_stbu**, EA is placed into GPR(rA).
- If **e\_stbu** and rA = 0, the instruction form is invalid.
- None



if (RA=0 & lse\_sth) then a ← <sup>32</sup>0 else a ← GPR(RA or RX)  
 if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>  
 if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>  
 if SD4-mode then EA ← (a + (<sup>27</sup>0 || SD4 || 0))<sub>32:63</sub>  
 MEM(EA,2) ← GPR(RS or RZ)<sub>48:63</sub>  
 if e\_sthu then GPR(RA) ← EA

Let the EA be calculated as follows:

- For **e\_sth** and **e\_sthu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_sth** let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 1 bit.

The contents of bits 48–63 of GPR(rS) are stored into the half word in memory addressed by EA.

If **e\_sthu**, EA is placed into GPR(rA).


If **e\_sthu** and rA = 0, the instruction form is invalid.

Special Registers Altered: None

**\_stmw** VLE User **\_stmw**

Store Multiple Word

**e\_stmw** **rS,D8(rA)** (D8-mode)



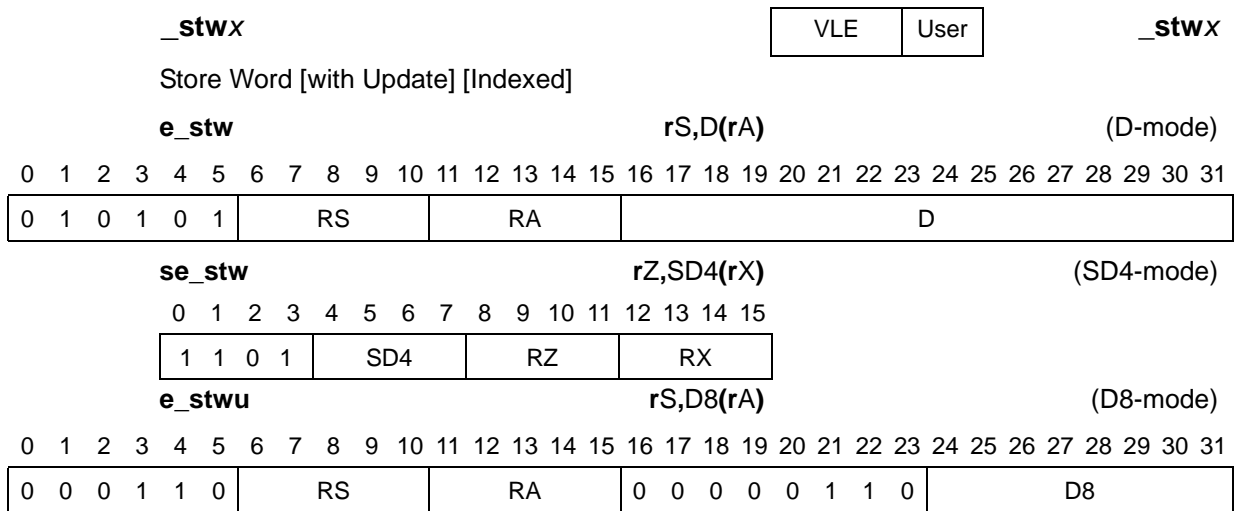
if RA=0 then EA ← EXTS(D8)<sub>32:63</sub>  
 else EA ← (GPR(RA)+EXTS(D8))<sub>32:63</sub>  
 r ← RS  
 do while r ≤ 31  
     MEM(EA,4) ← GPR(r)<sub>32:63</sub>  
     r ← r + 1  
     EA ← (EA+4)<sub>32:63</sub>

Let the EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D8 instruction field.

Let n = (32 - rS). Bits 32–63 of registers GPR(rS) through GPR(31) are stored in n consecutive words in memory starting at address EA.

EA must be a multiple of 4. If it is not, either an alignment interrupt is invoked or the results are boundedly undefined.

Special Registers Altered: None



if (RA=0 & lse\_stw) then a ← <sup>32</sup>0 else a ← GPR(RA or RX)  
 if D-mode then EA ← (a + EXTS(D))<sub>32:63</sub>  
 if D8-mode then EA ← (a + EXTS(D8))<sub>32:63</sub>  
 if SD4-mode then EA ← (a + (<sup>26</sup>0 || SD4 || <sup>2</sup>0))<sub>32:63</sub>  
 MEM(EA,4) ← GPR(RS or RZ)<sub>32:63</sub>

Let the EA be calculated as follows:


- For **e\_stw** and **e\_stwu**, let EA be the sum of the contents of GPR(rA), or 32 0s if rA = 0, and the sign-extended value of the D or D8 instruction field.
- For **se\_stw**, let EA be the sum of the contents of GPR(rX) and the zero-extended value of the SD4 instruction field shifted left by 2 bits.

The contents of bits 32–63 of GPR(rS) are stored into the word in memory addressed by EA.

If **e\_stwu**, EA is placed into GPR(rA).

If **e\_stwu** and rA = 0, the instruction form is invalid.

Special Registers Altered: None



sum<sub>32:63</sub> ← GPR(RX) + ¬GPR(RY) + 1  
 GPR(RX) ← sum<sub>32:63</sub>

The sum of the contents of GPR(rX), the one's complement of contents of GPR(rY), and 1 is placed into GPR(rX).


Special Registers Altered: None

**\_subfx**

VLE	User
-----	------

**\_subfx**

Subtract From



$$\text{sum}_{32:63} \leftarrow \neg\text{GPR}(\text{RX}) + \text{GPR}(\text{RY}) + 1$$

$$\text{GPR}(\text{RX}) \leftarrow \text{sum}_{32:63}$$

The sum of the one's complement of the contents of GPR(rX), the contents of GPR(rY), and 1 is placed into GPR(rX).

Special Registers Altered: None


**\_subfidx**

VLE	User
-----	------

**\_subfidx**

Subtract From Immediate Carrying [and Record]

**e\_subfic** **rD,rA,SCI8** (Rc = 0)  
**e\_subfic.** **rD,rA,SCI8** (Rc = 1)



$$\text{imm} \leftarrow \text{SCI8}(\text{F}, \text{SCL}, \text{UI8})$$

$$\text{carry}_{32:63} \leftarrow \text{Carry}(\neg\text{GPR}(\text{RA}) + \text{imm} + 1)$$

$$\text{sum}_{32:63} \leftarrow \neg\text{GPR}(\text{RA}) + \text{imm} + 1$$

if Rc=1 then do

$$\text{LT} \leftarrow \text{sum}_{32:63} < 0$$

$$\text{GT} \leftarrow \text{sum}_{32:63} > 0$$

$$\text{EQ} \leftarrow \text{sum}_{32:63} = 0$$

$$\text{CR0} \leftarrow \text{LT} \parallel \text{GT} \parallel \text{EQ} \parallel \text{SO}$$

$$\text{GPR}(\text{RD}) \leftarrow \text{sum}_{32:63}$$

$$\text{CA} \leftarrow \text{carry}_{32}$$

The sum of the one's complement of the contents of GPR(rA), the value of SCI8, and 1 is placed into GPR(rD).

Special Registers Altered: CA CR0 (if Rc=1)


**\_subix**

VLE	User
-----	------

**\_subix**

Subtract Immediate [and Record]

**se\_subi** **rX,OIMM** (Rc = 0)  
**se\_subi.** **rX,OIMM** (Rc = 1)



1. OIMM = OIM5 + 1



```

sum32:63 ← GPR(RX) + -(270 || OFFSET(OIM5)) + 1
if Rc=1 then do
    LT ← sum32:63 < 0
    GT ← sum32:63 > 0
    EQ ← sum32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RX) ← sum32:63
    
```

The sum of the contents of GPR(rX), the one's complement of the zero-extended value of the offseted OIM5 field (a final value in the range 1–32), and 1 is placed into GPR(rX).

Special Registers Altered: CR0 (if Rc = 1)

```

    _xorx                                     _xorx
XOR [Immediate] [and Record]
e_xori           rA,rS,SCI8                (Rc = 0)
e_xori.          rA,rS,SCI8                (Rc = 1)
    
```

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	0	1	1	0	RS						RA						1	1	1	0	Rc	F	SCL	UI8						

```

if 'e_xori[.]' then b ← SCI8(F,SCL,UI8)
result32:63 ← GPR(RS) ⊕ b
if Rc=1 then do
    LT ← result32:63 < 0
    GT ← result32:63 > 0
    EQ ← result32:63 = 0
    CR0 ← LT || GT || EQ || SO
GPR(RA) ← result
    
```

For **e\_xori[.]**, the contents of GPR(rS) are XORed with SCI8.

The result is placed into GPR(rA).

Special Registers Altered: CR0 (if Rc = 1)

# 15 VLE instruction index

The tables in this appendix use the following conventions:

**Table 260. Notation conventions**

Notation	Meaning
-	Don't care, usually part of an operand field
/	Reserved bit, invalid instruction form if encoded as 1
?	Allocated for implementation-dependent use. See the implementation documentation.

## 15.1 Instruction index sorted by opcode

*Table 261* lists the 16-bit VLE instructions, sorted by opcode.

**Table 261. Instruction index sorted by opcode**

Format	16-Bit opcodes	Mnemonic	Instruction	Page
	(Inst <sub>0:15</sub> )			
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	<b>se_illegal</b>	Illegal	<i>on page 757</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1	<b>se_isync</b>	Instruction Synchronize	<i>on page 757</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0	<b>se_sc</b>	System Call	<i>on page 769</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0	<b>se_blr</b>	Branch to Link Register	<i>on page 747</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1	<b>se_blrl</b>	Branch to link register & link	<i>on page 747</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 0	<b>se_bctr</b>	Branch to Count Register	<i>on page 746</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1	<b>se_bctrl</b>	Branch to Count Register & Link	<i>on page 746</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0	<b>se_rfi</b>	Return From Interrupt	<i>on page 767</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1	<b>se_rfci</b>	Return From Critical Interrupt	<i>on page 767</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1	<b>se_rfdi</b>	Return From Debug Interrupt	<i>on page 767</i>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1			
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 - - - -			
R	0 0 0 0 0 0 0 0 0 0 0 0 1 0 x x x x	<b>se_not</b>	NOT	<i>on page 764</i>

Table 261. Instruction index sorted by opcode (continued)

Format	16-Bit opcodes	Mnemonic	Instruction	Page
	(Inst <sub>0:15</sub> )			
R	0 0 0 0 0 0 0 0 1 1 x x x x	<b>se_neg</b>	Negate	<a href="#">on page 764</a>
R	0 0 0 0 0 0 0 0 1 - - x x x x			
R	0 0 0 0 0 0 0 0 1 0 0 0 x x x x	<b>se_mflr</b>	Move From Link Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 0 1 x x x x	<b>se_mtlr</b>	Move To Link Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 0 1 0 x x x x	<b>se_mfctr</b>	Move From Count Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 1 1 x x x x	<b>se_mtctr</b>	Move To Count Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 1 0 0 x x x x	<b>se_extzb</b>	Extend with Zeros Byte	<a href="#">on page 756</a>
R	0 0 0 0 0 0 0 0 1 1 0 1 x x x x	<b>se_extsb</b>	Extend Sign Byte	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 0 1 1 1 0 x x x x	<b>se_extzh</b>	Extend with Zeros Halfword	<a href="#">on page 756</a>
R	0 0 0 0 0 0 0 0 1 1 1 1 x x x x	<b>se_extsh</b>	Extend Sign Halfword	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 1 y y y y x x x x	<b>se_mr</b>	Move Register	<a href="#">on page 762</a>
RR	0 0 0 0 0 0 1 0 y y y y x x x x	<b>se_mtar</b>	Move to Alternate Register	<a href="#">on page 763</a>
RR	0 0 0 0 0 0 1 1 y y y y x x x x	<b>se_mfar</b>	Move from Alternate Register	<a href="#">on page 762</a>
RR	0 0 0 0 0 1 0 0 y y y y x x x x	<b>se_add</b>	Add	<a href="#">on page 741</a>
RR	0 0 0 0 0 1 0 1 y y y y x x x x	<b>se_mullw</b>	Multiply Low Word	<a href="#">on page 764</a>
RR	0 0 0 0 0 1 1 0 y y y y x x x x	<b>se_sub</b>	Subtract	<a href="#">on page 775</a>
RR	0 0 0 0 0 1 1 1 y y y y x x x x	<b>se_subf</b>	Subtract From	<a href="#">on page 776</a>
RR	0 0 0 0 1 0 - - y y y y x x x x			
RR	0 0 0 0 1 1 0 0 y y y y x x x x	<b>se_cmp</b>	Compare	<a href="#">on page 749</a>
RR	0 0 0 0 1 1 0 1 y y y y x x x x	<b>se_cmpl</b>	Compare Logical	<a href="#">on page 752</a>
RR	0 0 0 0 1 1 1 0 y y y y x x x x	<b>se_cmph</b>	Compare Halfword	<a href="#">on page 750</a>

Table 261. Instruction index sorted by opcode (continued)

Format	16-Bit opcodes				Mnemonic	Instruction	Page												
	(Inst <sub>0:15</sub> )																		
RR	0	0	0	0	1	1	1	1	y	y	y	y	x	x	x	x	<b>se_cmphl</b>	Compare Halfword Logical	<i>on page 751</i>
IM5	0	0	1	0	0	0	0	0	i	i	i	i	x	x	x	x	<b>se_addi</b>	Add Immediate	<i>on page 741</i>
IM5	0	0	1	0	0	0	1	i	i	i	i	i	x	x	x	x	<b>se_cmpli</b>	Compare Logical Immediate	<i>on page 752</i>
IM5	0	0	1	0	0	1	0	i	i	i	i	i	x	x	x	x	<b>se_subi</b>	Subtract Immediate	<i>on page 776</i>
IM5	0	0	1	0	0	1	1	i	i	i	i	i	x	x	x	x	<b>se_subi.</b>	Subtract Immediate and Record	<i>on page 776</i>
IM5	0	0	1	0	1	0	0	i	i	i	i	i	x	x	x	x			
IM5	0	0	1	0	1	0	1	i	i	i	i	i	x	x	x	x	<b>se_cmpi</b>	Compare Immediate	<i>on page 749</i>
IM5	0	0	1	0	1	1	0	i	i	i	i	i	x	x	x	x	<b>se_bmask i</b>	Bit Mask Generate Immediate	<i>on page 748</i>
IM5	0	0	1	0	1	1	1	i	i	i	i	i	x	x	x	x	<b>se_andi</b>	And Immediate	<i>on page 743</i>
RR	0	1	0	0	0	0	0	0	y	y	y	y	x	x	x	x	<b>se_srw</b>	Shift Right Word	<i>on page 772</i>
RR	0	1	0	0	0	0	0	1	y	y	y	y	x	x	x	x	<b>se_sraw</b>	Shift Right Algebraic Word	<i>on page 771</i>
RR	0	1	0	0	0	0	1	0	y	y	y	y	x	x	x	x	<b>se_slw</b>	Shift Left Word	<i>on page 770</i>
RR	0	1	0	0	0	0	1	1	y	y	y	y	x	x	x	x			
RR	0	1	0	0	0	1	0	0	y	y	y	y	x	x	x	x	<b>se_or</b>	OR	<i>on page 765</i>
RR	0	1	0	0	0	1	0	1	y	y	y	y	x	x	x	x	<b>se_andc</b>	AND with Complement	<i>on page 743</i>
RR	0	1	0	0	0	1	1	0	y	y	y	y	x	x	x	x	<b>se_and</b>	AND	<i>on page 743</i>
RR	0	1	0	0	0	1	1	1	y	y	y	y	x	x	x	x	<b>se_and.</b>	AND and Record	<i>on page 743</i>
IM7	0	1	0	0	1	i	i	i	i	i	i	i	x	x	x	x	<b>se_li</b>	Load Immediate	<i>on page 760</i>
IM5	0	1	1	0	0	0	0	i	i	i	i	i	x	x	x	x	<b>se_bclri</b>	Bit Clear Immediate	<i>on page 746</i>
IM5	0	1	1	0	0	0	1	i	i	i	i	i	x	x	x	x	<b>se_bgeni</b>	Bit Generate Immediate	<i>on page 746</i>
IM5	0	1	1	0	0	1	0	i	i	i	i	i	x	x	x	x	<b>se_bseti</b>	Bit Set Immediate	<i>on page 748</i>

Table 261. Instruction index sorted by opcode (continued)

Format	16-Bit opcodes	Mnemonic	Instruction	Page
	(Inst <sub>0:15</sub> )			
IM5	0 1 1 0 0 1 1 i i i i x x x x	<b>se_btsti</b>	Bit Test Immediate	<a href="#">on page 748</a>
IM5	0 1 1 0 1 0 0 i i i i x x x x	<b>se_srwi</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
IM5	0 1 1 0 1 0 1 i i i i x x x x	<b>se_srawi</b>	Shift Right Algebraic Word Immediate	<a href="#">on page 772</a>
IM5	0 1 1 0 1 1 0 i i i i x x x x	<b>se_slwi</b>	Shift Left Word Immediate	<a href="#">on page 770</a>
IM5	0 1 1 0 1 1 1 i i i i x x x x			
SD4	1 0 0 0 i i i i z z z z x x x x	<b>se_lbz</b>	Load Byte and Zero	<a href="#">on page 757</a>
SD4	1 0 0 1 i i i i z z z z x x x x	<b>se_stb</b>	Store Byte	<a href="#">on page 773</a>
SD4	1 0 1 0 i i i i z z z z x x x x	<b>se_lhz</b>	Load Halfword and Zero	<a href="#">on page 759</a>
SD4	1 0 1 1 i i i i z z z z x x x x	<b>se_sth</b>	Store Halfword	<a href="#">on page 773</a>
SD4	1 1 0 0 i i i i z z z z x x x x	<b>se_lwz</b>	Load Word and Zero	<a href="#">on page 761</a>
SD4	1 1 0 1 i i i i z z z z x x x x	<b>se_stw</b>	Store Word	<a href="#">on page 775</a>
B8	1 1 1 0 0 o i i d d d d d d d d	<b>se_bc</b>	Branch Conditional	<a href="#">on page 745</a>
B8	1 1 1 0 1 0 0 0 d d d d d d d d	<b>se_b</b>	Branch	<a href="#">on page 748</a>
B8	1 1 1 0 1 0 0 1 d d d d d d d d	<b>se_bl</b>	Branch and Link	<a href="#">on page 748</a>
	1 1 1 0 1 0 1 - - - - - - - -			
	1 1 1 0 1 1 - - - - - - - -			

Table 262 shows 32-bit instruction encodings.

Table 262. 32-bit instruction encodings

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
APU	0 0 0 1 0 -	- - - - -	- - - - -	<b>apu</b>	Reserved for APUs	
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	0 0 0 d d d d d d d d	<b>e_lbzu</b>	Load Byte & Zero with Update	<a href="#">on page 757</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	0 0 1 d d d d d d d d	<b>e_lhzu</b>	Load Halfword & Zero with Update	<a href="#">on page 759</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	0 1 0 d d d d d d d d	<b>e_lwzu</b>	Load Word & Zero with Update	<a href="#">on page 761</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	0 1 1 d d d d d d d d	<b>e_lhau</b>	Load Halfword Algebraic With Update	<a href="#">on page 758</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 0 0 d d d d d d d d	<b>e_stbu</b>	Store Byte with Update	<a href="#">on page 773</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 0 1 d d d d d d d d	<b>e_sthu</b>	Store Halfword with Update	<a href="#">on page 773</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 1 0 d d d d d d d d	<b>e_stwu</b>	Store Word with Update	<a href="#">on page 775</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 1 1 d d d d d d d d			
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 1	0 0 0 d d d d d d d d	<b>e_lmw</b>	Load Multiple Word	<a href="#">on page 760</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 1	0 0 1 d d d d d d d d	<b>e_stmw</b>	Store Multiple Word	<a href="#">on page 774</a>
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 1	0 1 0 d d d d d d d d			
D8	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 1	0 1 1 d d d d d d d d			
	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 0 1	1 - - d d d d d d d d			
	0 0 0 1 1 0	t t t t t a a a a a 0 0 0 1 -	- - - - -			
	0 0 0 1 1 0	t t t t t a a a a a 0 0 1 - -	- - - - -			
	0 0 0 1 1 0	t t t t t a a a a a 0 1 - - -	- - - - -			
SC18	0 0 0 1 1 0	t t t t t a a a a a 1 0 0 0 0	F S S i i i i i i i i	<b>e_addi</b>	Add Immediate	<a href="#">on page 741</a>



Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
SCI8	000110	tttttt aaaaa 10001	FSSi i i i i i i i	<b>e_addi.</b>	Add Immediate and Record	<a href="#">on page 741</a>
SCI8	000110	tttttt aaaaa 10010	FSSi i i i i i i i	<b>e_addic</b>	Add Immediate Carrying	<a href="#">on page 743</a>
SCI8	000110	tttttt aaaaa 10011	FSSi i i i i i i i	<b>e_addic.</b>	Add Immediate Carrying and Record	<a href="#">on page 743</a>
SCI8	000110	tttttt aaaaa 10100	FSSi i i i i i i i	<b>e_mulli</b>	Multiply Low Immediate	<a href="#">on page 763</a>
SCI8	000110	000bf aaaaa 10101	FSSi i i i i i i i	<b>e_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
SCI8	000110	001bf aaaaa 10101	FSSi i i i i i i i	<b>e_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
SCI8	000110	tttttt aaaaa 10110	FSSi i i i i i i i	<b>e_subfic</b>	Subtract from Immediate Carrying	<a href="#">on page 776</a>
SCI8	000110	tttttt aaaaa 10111	FSSi i i i i i i i	<b>e_subfic.</b>	Subtract from Immediate and Record	<a href="#">on page 776</a>
SCI8	000110	sssss aaaaa 11000	FSSi i i i i i i i	<b>e_andi</b>	AND Immediate	<a href="#">on page 743</a>
SCI8	000110	sssss aaaaa 11001	FSSi i i i i i i i	<b>e_andi.</b>	AND Immediate and Record	<a href="#">on page 743</a>
SCI8	000110	sssss aaaaa 11010	FSSi i i i i i i i	<b>e_ori</b>	OR Immediate	<a href="#">on page 768</a>
SCI8	000110	sssss aaaaa 11011	FSSi i i i i i i i	<b>e_ori.</b>	OR Immediate and Record	<a href="#">on page 768</a>
SCI8	000110	sssss aaaaa 11100	FSSi i i i i i i i	<b>e_xori</b>	XOR Immediate	<a href="#">on page 777</a>
SCI8	000110	sssss aaaaa 11101	FSSi i i i i i i i	<b>e_xori.</b>	XOR Immediate and Record	<a href="#">on page 777</a>
SCI8	000110	sssss aaaaa 11110	FSSi i i i i i i i			



Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
SCI8	000110	sssssa aaaaa 11111	FSSi i i i i i i i			
D	000111	ttttt aaaaa i i i i i	i i i i i i i i i i	<b>e_add16i</b>	Add Immediate	<a href="#">on page 741</a>
D	001100	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_lbz</b>	Load Byte & Zero	<a href="#">on page 757</a>
D	001101	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_stb</b>	Store Byte	<a href="#">on page 773</a>
D	001110	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_lha</b>	Load Halfword Algebraic	<a href="#">on page 758</a>
	001111	- - - - -	- - - - -			
D	010100	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_lwz</b>	Load Word & Zero	<a href="#">on page 761</a>
D	010101	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_stw</b>	Store Word	<a href="#">on page 775</a>
D	010110	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_lhz</b>	Load Halfword & Zero	<a href="#">on page 759</a>
D	010111	ttttt aaaaa d d d d d	d d d d d d d d d d	<b>e_sth</b>	Store Halfword	<a href="#">on page 773</a>
LI20	011100	ttttt i i i i i 0 i i i i	i i i i i i i i i i	<b>e_li</b>	Load Immediate	<a href="#">on page 760</a>
I16A	011100	i i i i i aaaaa 10000	i i i i i i i i i i			
I16A	011100	i i i i i aaaaa 10001	i i i i i i i i i i	<b>e_add2i.</b>	Add (2 operand) Immediate and Record CR	<a href="#">on page 741</a>
I16A	011100	i i i i i aaaaa 10010	i i i i i i i i i i	<b>e_add2is</b>	Add (2 operand) Immediate Shifted	<a href="#">on page 741</a>
I16A	011100	i i i i i aaaaa 10011	i i i i i i i i i i	<b>e_cmp16i</b>	Compare Immediate	<a href="#">on page 749</a>





Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
I16A	011100	iiiiiaaaaa	10100iiiiiiii	<b>e_mull2i</b>	Multiply Low Word (2 operand) Immediate	<a href="#">on page 763</a>
I16A	011100	iiiiiaaaaa	10101iiiiiiii	<b>e_cmpl16i</b>	Compare Logical Immediate	<a href="#">on page 752</a>
I16A	011100	iiiiiaaaaa	10110iiiiiiii	<b>e_cmph16i</b>	Compare Halfword Immediate	<a href="#">on page 750</a>
I16A	011100	iiiiiaaaaa	10111iiiiiiii	<b>e_cmph16i</b>	Compare Halfword Logical Immediate	<a href="#">on page 751</a>
I16L	011100	tttttiiii	11000iiiiiiii	<b>e_or2i</b>	OR (2 operand) Immediate	<a href="#">on page 765</a>
I16L	011100	tttttiiii	11001iiiiiiii	<b>e_and2i.</b>	AND (2 operand) Immediate & record CR	<a href="#">on page 743</a>
I16L	011100	tttttiiii	11010iiiiiiii	<b>e_or2is</b>	OR (2 operand) Immediate Shifted	<a href="#">on page 777</a>
I16L	011100	tttttiiii	11011iiiiiiii			
I16L	011100	tttttiiii	11100iiiiiiii	<b>e_lis</b>	Load Immediate Shifted	<a href="#">on page 760</a>
I16L	011100	tttttiiii	11101iiiiiiii	<b>e_and2is.</b>	AND (2 operand) Immediate Shifted & record CR	<a href="#">on page 743</a>
I16L	011100	tttttiiii	11110iiiiiiii			
I16L	011100	tttttiiii	11111iiiiiiii			
RLWI	011101	ssssaaaa	hhhhbbbbbb	<b>e_rlwimi</b>	Rotate Left Word Immed then Mask Insert	<a href="#">on page 768</a>
RLWI	011101	ssssaaaa	hhhhbbbbbb	<b>e_rlwinm</b>	Rotate Left Word Immed then AND with Mask	<a href="#">on page 769</a>



Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
BD24	0 1 1 1 1 0	0 d d d d d d d d d d d d d d d d	d d d d d d d d d d d d d d d d 0	<b>e_b</b>	Branch	<a href="#">on page 744</a>
BD24	0 1 1 1 1 0	0 d d d d d d d d d d d d d d d d	d d d d d d d d d d d d d d d d 1	<b>e_bl</b>	Branch & Link	<a href="#">on page 744</a>
BD15	0 1 1 1 1 0	1 0 0 0 0 o o i i i i d d d d d d	d d d d d d d d d d d d d d d d 0	<b>e_bc</b>	Branch Conditional	<a href="#">on page 745</a>
BD15	0 1 1 1 1 0	1 0 0 0 0 o o i i i i d d d d d d	d d d d d d d d d d d d d d d d 1	<b>e_bcl</b>	Branch Conditional & Link	<a href="#">on page 745</a>
X	0 1 1 1 1 1		- - - - - 0 1 1 1 1 /	<b>isel</b>	Integer Select	Book E
X	0 1 1 1 1 1		/ 0 0 0 0 0 1 0 1 1 0	<b>mulhwu</b>	Multiply High Word Unsigned	Book E
X	0 1 1 1 1 1		/ 0 0 0 0 0 1 0 1 1 1	<b>mulhwu.</b>	Multiply High Word Unsigned & Record	Book E
X	0 1 1 1 1 1		/ 0 0 1 0 0 1 0 1 1 0	<b>mulhw</b>	Multiply High Word	Book E
X	0 1 1 1 1 1		/ 0 0 1 0 0 1 0 1 1 1	<b>mulhw.</b>	Multiply High Word & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 0 0 0 0 /	<b>cmp</b>	Compare	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 0 1 0 0 /	<b>tw</b>	Trap Word	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 1 0 0 0 0	<b>subfc</b>	Subtract From Carrying	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 1 0 0 0 1	<b>subfc.</b>	Subtract From Carrying & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 1 0 1 0 0	<b>addc</b>	Add Carrying	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 1 0 1 0 1	<b>addc.</b>	Add Carrying & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 0 0 1 1 1 0 /	<b>e_cmph</b>	Compare Halfword	<a href="#">on page 750</a>



Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
XL	0 1 1 1 1 1		0 0 0 0 0 1 0 0 0 0 /	<b>e_mcrf</b>	Move Condition Register Field	<a href="#">on page 763</a>
X	0 1 1 1 1 1		0 0 0 0 0 1 0 0 1 1 /	<b>mfcrr</b>	Move From Condition Register	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 0 1 0 0 /	<b>lwarx</b>	Load Word & Reserve Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 0 1 1 0 /	<b>icbt</b>	Instruction Cache Block Touch Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 0 1 1 1 /	<b>lwzxx</b>	Load Word & Zero Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 0 0 0	<b>slw</b>	Shift Left Word	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 0 0 1	<b>slw.</b>	Shift Left Word & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 1 0 0	<b>cntlzw</b>	Count Leading Zeros Word	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 1 0 1	<b>cntlzw.</b>	Count Leading Zeros Word & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 1 0 0 0	<b>and</b>	AND	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 1 0 0 1	<b>and.</b>	AND & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 1 0 0 0 0 0 /	<b>cmpl</b>	Compare Logical	Book E
XL	0 1 1 1 1 1		0 0 0 0 1 0 0 0 0 1 /	<b>e_crnor</b>	Condition Register NOR	<a href="#">on page 754</a>
X	0 1 1 1 1 1		0 0 0 0 1 0 1 0 0 0 0	<b>subf</b>	Subtract From	Book E
X	0 1 1 1 1 1		0 0 0 0 1 0 1 0 0 0 1	<b>subf.</b>	Subtract From & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 1 0 1 1 1 0 /	<b>e_cmphi</b>	Compare Halfword Logical	<a href="#">on page 751</a>

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 0 0 0 1 1 0 1 1 0 /	<b>dcbst</b>	Data Cache Block Store Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 1 1 0 1 1 1 /	<b>lwzux</b>	Load Word & Zero with Update Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 1 1 1 0 0 0 0	<b>e_slwi</b>	Shift Left Word Immediate	<a href="#">on page 770</a>
X	0 1 1 1 1 1		0 0 0 0 1 1 1 0 0 0 1	<b>e_slwi.</b>	Shift Left Word Immediate & record CR	<a href="#">on page 770</a>
X	0 1 1 1 1 1		0 0 0 0 1 1 1 1 0 0 0	<b>andc</b>	AND with Complement	Book E
X	0 1 1 1 1 1		0 0 0 0 1 1 1 1 0 0 1	<b>andc.</b>	AND with Complement & record CR	Book E
X	0 1 1 1 1 1		0 0 0 1 0 1 0 0 1 1 /	<b>mfmsr</b>	Move From Machine State Register	Book E
X	0 1 1 1 1 1		0 0 0 1 0 1 0 1 1 0 /	<b>dcbf</b>	Data Cache Block Flush Indexed	Book E
X	0 1 1 1 1 1		0 0 0 1 0 1 0 1 1 1 /	<b>lbzx</b>	Load Byte & Zero Indexed	Book E
X	0 1 1 1 1 1		0 0 0 1 1 0 1 0 0 0 0	<b>neg</b>	Negate	Book E
X	0 1 1 1 1 1		0 0 0 1 1 0 1 0 0 0 1	<b>neg.</b>	Negate & record CR	Book E
X	0 1 1 1 1 1		0 0 0 1 1 1 0 1 1 1 /	<b>lbzux</b>	Load Byte & Zero with Update Indexed	Book E
X	0 1 1 1 1 1		0 0 0 1 1 1 1 1 0 0 0	<b>nor</b>	NOR	Book E
X	0 1 1 1 1 1		0 0 0 1 1 1 1 1 0 0 1	<b>nor.</b>	NOR & record CR	Book E
XL	0 1 1 1 1 1		0 0 1 0 0 0 0 0 0 1 /	<b>e_crandc</b>	Condition Register AND with Complement	<a href="#">on page 753</a>
X	0 1 1 1 1 1		0 0 1 0 0 0 0 0 1 1 /	<b>wrtee</b>	Write External Enable	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 0 1 0 0 0 1 0 0 0 0	<b>subfe</b>	Subtract From Extended with CA	Book E
X	0 1 1 1 1 1		0 0 1 0 0 0 1 0 0 0 1	<b>subfe.</b>	Subtract From Extended with CA & record CR	Book E
X	0 1 1 1 1 1		0 0 1 0 0 0 1 0 1 0 0	<b>adde</b>	Add Extended with CA	Book E
X	0 1 1 1 1 1		0 0 1 0 0 0 1 0 1 0 1	<b>adde.</b>	Add Extended with CA & record CR	Book E
AFX	0 1 1 1 1 1		0 0 1 0 0 1 0 0 0 0 /	<b>mtrcf</b>	Move To Condition Register Fields	Book E
X	0 1 1 1 1 1		0 0 1 0 0 1 0 0 1 0 /	<b>mtmsr</b>	Move To Machine State Register	Book E
X	0 1 1 1 1 1		0 0 1 0 0 1 0 1 1 0 1	<b>stwcx.</b>	Store Word Conditional Indexed & record CR	Book E
X	0 1 1 1 1 1		0 0 1 0 0 1 0 1 1 1 /	<b>stwx</b>	Store Word Indexed	Book E
X	0 1 1 1 1 1		0 0 1 0 1 0 0 0 1 1 /	<b>wrteei</b>	Write External Enable Immediate	Book E
X	0 1 1 1 1 1		0 0 1 0 1 1 0 1 1 1 /	<b>stwux</b>	Store Word with Update Indexed	Book E
XL	0 1 1 1 1 1		0 0 1 1 0 0 0 0 0 1 /	<b>e_crxor</b>	Condition Register XOR	<i>on page 755</i>
X	0 1 1 1 1 1		0 0 1 1 0 0 1 0 0 0 0	<b>subfze</b>	Subtract From Zero Extended with CA	Book E
X	0 1 1 1 1 1		0 0 1 1 0 0 1 0 0 0 1	<b>subfze.</b>	Subtract From Zero Extended with CA & record CR	Book E
X	0 1 1 1 1 1		0 0 1 1 0 0 1 0 1 0 0	<b>addze</b>	Add to Zero Extended with CA	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 0 1 1 0 0 1 0 1 0 1	<b>addze.</b>	Add to Zero Extended with CA & record CR	Book E
X	0 1 1 1 1 1		0 0 1 1 0 1 0 1 1 1 /	<b>stbx</b>	Store Byte Indexed	Book E
XL	0 1 1 1 1 1		0 0 1 1 1 0 0 0 0 1 /	<b>e_crnand</b>	Condition Register NAND	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 0 0 0	<b>subfme</b>	Subtract From Minus One Extended with CA	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 0 0 1	<b>subfme.</b>	Subtract From Minus One Extended with CA & record CR	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 1 0 0	<b>addme</b>	Add to Minus One Extended with CA	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 1 0 1	<b>addme.</b>	Add to Minus One Extended with CA & record CR	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 1 1 0	<b>mullw</b>	Multiply Low Word	Book E
X	0 1 1 1 1 1		0 0 1 1 1 0 1 0 1 1 1	<b>mullw.</b>	Multiply Low Word & record CR	Book E
X	0 1 1 1 1 1		0 0 1 1 1 1 0 1 1 1 0 /	<b>dcbtst</b>	Data Cache Block Touch for Store Indexed	Book E
X	0 1 1 1 1 1		0 0 1 1 1 1 0 1 1 1 1 /	<b>stbux</b>	Store Byte with Update Indexed	Book E
XL	0 1 1 1 1 1		0 1 0 0 0 0 0 0 0 1 /	<b>e_crand</b>	Condition Register AND	<i>on page 753</i>
X	0 1 1 1 1 1		0 1 0 0 0 0 1 0 1 0 0	<b>add</b>	Add	Book E
X	0 1 1 1 1 1		0 1 0 0 0 0 1 0 1 0 1	<b>add.</b>	Add & record CR	Book E
X	0 1 1 1 1 1		0 1 0 0 0 1 0 0 1 1 /	<b>mfapidi</b>	Move From APID Indirect	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 1 0 0 0 1 0 1 1 0 /	<b>dcbt</b>	Data Cache Block Touch Indexed	Book E
X	0 1 1 1 1 1		0 1 0 0 0 1 0 1 1 1 /	<b>lhzx</b>	Load Halfword & Zero Indexed	Book E
X	0 1 1 1 1 1		0 1 0 0 0 1 1 0 0 0 0	<b>e_rlw</b>	Rotate Left Word	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 0 1 1 0 0 0 1	<b>e_rlw.</b>	Rotate Left Word & record CR	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 0 1 1 1 0 0 0	<b>eqv</b>	Equivalent	Book E
X	0 1 1 1 1 1		0 1 0 0 0 1 1 1 0 0 1	<b>eqv.</b>	Equivalent & record CR	Book E
XL	0 1 1 1 1 1		0 1 0 0 1 0 0 0 0 1 /	<b>e_creqv</b>	Condition Register Equivalent	<a href="#">on page 753</a>
X	0 1 1 1 1 1		0 1 0 0 1 1 0 1 1 1 /	<b>lhzux</b>	Load Halfword & Zero with Update Indexed	Book E
X	0 1 1 1 1 1		0 1 0 0 1 1 1 0 0 0 0	<b>e_rlwi</b>	Rotate Left Word Immediate	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 1 1 1 0 0 0 1	<b>e_rlwi.</b>	Rotate Left Word Immediate & record CR	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 1 1 1 1 0 0 0	<b>xor</b>	XOR	Book E
X	0 1 1 1 1 1		0 1 0 0 1 1 1 1 0 0 1	<b>xor.</b>	XOR & record CR	Book E
XFX	0 1 1 1 1 1		0 1 0 1 0 0 0 0 1 1 /	<b>mf dcr</b>	Move From Device Control Register	Book E
XFX	0 1 1 1 1 1		0 1 0 1 0 1 0 0 1 1 /	<b>mf spr</b>	Move From Special Purpose Register	Book E
X	0 1 1 1 1 1		0 1 0 1 0 1 0 1 1 1 /	<b>lhax</b>	Load Halfword Algebraic Indexed	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 1 0 1 1 1 0 1 1 1 1 /	<b>lhax</b>	Load Halfword Algebraic with Update Indexed	Book E
X	0 1 1 1 1 1		0 1 1 0 0 1 0 1 1 1 1 /	<b>sthx</b>	Store Halfword Indexed	Book E
X	0 1 1 1 1 1		0 1 1 0 0 1 1 1 0 0 0	<b>orc</b>	OR with Complement	Book E
X	0 1 1 1 1 1		0 1 1 0 0 1 1 1 0 0 1	<b>orc.</b>	OR with Complement & record CR	Book E
XL	0 1 1 1 1 1		0 1 1 0 1 0 0 0 0 1 /	<b>e_crorc</b>	Condition Register OR with Complement	<a href="#">on page 755</a>
X	0 1 1 1 1 1		0 1 1 0 1 1 0 1 1 1 1 /	<b>sthux</b>	Store Halfword with Update Indexed	Book E
X	0 1 1 1 1 1		0 1 1 0 1 1 1 1 0 0 0	<b>or</b>	OR	Book E
X	0 1 1 1 1 1		0 1 1 0 1 1 1 1 0 0 1	<b>or.</b>	OR & record CR	Book E
XL	0 1 1 1 1 1		0 1 1 1 0 0 0 0 0 1 /	<b>e_cror</b>	Condition Register OR	<a href="#">on page 755</a>
AFX	0 1 1 1 1 1		0 1 1 1 0 0 0 0 1 1 /	<b>mtdcr</b>	Move To Device Control Register	Book E
X	0 1 1 1 1 1		0 1 1 1 0 0 1 0 1 1 0	<b>divwu</b>	Divide Word Unsigned	Book E
X	0 1 1 1 1 1		0 1 1 1 0 0 1 0 1 1 1	<b>divwu.</b>	Divide Word Unsigned & record CR	Book E
AFX	0 1 1 1 1 1		0 1 1 1 0 1 0 0 1 1 /	<b>mtspr</b>	Move To Special Purpose Register	Book E
X	0 1 1 1 1 1		0 1 1 1 0 1 0 1 1 0 /	<b>dcbi</b>	Data Cache Block Invalidate Indexed	Book E
X	0 1 1 1 1 1		0 1 1 1 0 1 1 1 0 0 0	<b>nand</b>	NAND	Book E



Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 1 1 1 0 1 1 1 0 0 1	<b>nand.</b>	NAND & record CR	Book E
X	0 1 1 1 1 1		0 1 1 1 1 0 1 0 1 1 0	<b>divw</b>	Divide Word	Book E
X	0 1 1 1 1 1		0 1 1 1 1 0 1 0 1 1 1	<b>divw.</b>	Divide Word & record CR	Book E
X	0 1 1 1 1 1		1 0 0 0 0 0 0 0 0 0 /	<b>mcrxr</b>	Move to Condition Register from XER	Book E
X	0 1 1 1 1 1		1 0 0 0 0 0 1 0 0 0 0	<b>subfco</b>	Subtract From Carrying & record OV	Book E
X	0 1 1 1 1 1		1 0 0 0 0 0 1 0 0 0 1	<b>subfco.</b>	Subtract From Carrying & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 0 0 0 0 1 0 1 0 0	<b>addco</b>	Add Carrying & record OV	Book E
X	0 1 1 1 1 1		1 0 0 0 0 0 1 0 1 0 1	<b>addco.</b>	Add Carrying & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 0 0 0 1 0 1 1 0 /	<b>lwbrx</b>	Load Word Byte-Reverse Indexed	Book E
X	0 1 1 1 1 1		1 0 0 0 0 1 1 0 0 0 0	<b>srw</b>	Shift Right Word	Book E
X	0 1 1 1 1 1		1 0 0 0 0 1 1 0 0 0 1	<b>srw.</b>	Shift Right Word & record CR	Book E
X	0 1 1 1 1 1		1 0 0 0 1 0 1 0 0 0 0	<b>subfo</b>	Subtract From & record OV	Book E
X	0 1 1 1 1 1		1 0 0 0 1 0 1 0 0 0 1	<b>subfo.</b>	Subtract From & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 0 0 1 1 0 1 1 0 /	<b>tlbsync</b>	TLB Synchronize	Book E
X	0 1 1 1 1 1		1 0 0 0 1 1 1 0 0 0 0	<b>e_srwi</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
X	0 1 1 1 1 1		1 0 0 0 1 1 1 0 0 0 1	<b>e_srwi.</b>	Shift Right Word Immediate & record CR	<a href="#">on page 772</a>

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		1 0 0 1 0 1 0 1 1 0 /	<b>msync</b>	Memory Synchronize	Book E
X	0 1 1 1 1 1		1 0 0 1 1 0 1 0 0 0 0	<b>nego</b>	Negate & record OV	Book E
X	0 1 1 1 1 1		1 0 0 1 1 0 1 0 0 0 1	<b>nego.</b>	Negate & record OV & record CR	Book E
X	0 1 1 1 1 1		1 0 1 0 0 0 1 0 0 0 0	<b>subfeo</b>	Subtract From Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 0 0 0 1 0 0 0 1	<b>subfeo.</b>	Subtract From Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 0 0 0 1 0 1 0 0	<b>addeo</b>	Add Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 0 0 0 1 0 1 0 1	<b>addeo.</b>	Add Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 0 0 1 0 1 0 1 /	<b>stswx</b>	Store String Word Indexed	Book E
X	0 1 1 1 1 1		1 0 1 0 0 1 0 1 1 0 /	<b>stwbrx</b>	Store Word Byte-Reverse Indexed	Book E
X	0 1 1 1 1 1		1 0 1 1 0 0 1 0 0 0 0	<b>subfzeo</b>	Subtract From Zero Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 1 0 0 1 0 0 0 1	<b>subfzeo.</b>	Subtract From Zero Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 1 0 0 1 0 1 0 0	<b>addzeo</b>	Add to Zero Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 1 0 0 1 0 1 0 1	<b>addzeo.</b>	Add to Zero Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 1 0 1 0 1 0 1 /	<b>stswi</b>	Store String Word Immediate	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 0 0 0	<b>subfmeo</b>	Subtract From Minus One Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 0 0 1	<b>subfmeo.</b>	Subtract From Minus One Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 1 0 0	<b>addmeo</b>	Add to Minus One Extended with CA & record OV	Book E
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 1 0 1	<b>addmeo.</b>	Add to Minus One Extended with CA & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 1 1 0	<b>mullwo</b>	Multiply Low Word & record OV	Book E
X	0 1 1 1 1 1		1 0 1 1 1 0 1 0 1 1 1	<b>mullwo.</b>	Multiply Low Word & record OV & CR	Book E
X	0 1 1 1 1 1		1 0 1 1 1 1 0 1 1 0 /	<b>dcba</b>	Data Cache Block Allocate Indexed	Book E
X	0 1 1 1 1 1		1 1 0 0 0 0 1 0 1 0 0	<b>addo</b>	Add & record OV	Book E
X	0 1 1 1 1 1		1 1 0 0 0 0 1 0 1 0 1	<b>addo.</b>	Add & record OV & CR	Book E
X	0 1 1 1 1 1		1 1 0 0 0 1 0 0 1 0 /	<b>tlbivax</b>	TLB Invalidate Virtual Address Indexed	Book E
X	0 1 1 1 1 1		1 1 0 0 0 1 0 1 1 0 /	<b>lhbrx</b>	Load Halfword Byte-Reverse Indexed	Book E
X	0 1 1 1 1 1		1 1 0 0 0 1 1 0 0 0 0	<b>sraw</b>	Shift Right Algebraic Word	Book E
X	0 1 1 1 1 1		1 1 0 0 0 1 1 0 0 0 1	<b>sraw.</b>	Shift Right Algebraic Word & record CR	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		1 1 0 0 1 1 1 0 0 0 0	<b>srawi</b>	Shift Right Algebraic Word Immediate	Book E
X	0 1 1 1 1 1		1 1 0 0 1 1 1 0 0 0 1	<b>srawi.</b>	Shift Right Algebraic Word Immediate & record CR	Book E
X	0 1 1 1 1 1		1 1 0 1 0 1 0 1 1 0 /	<b>mbar</b>	Memory Barrier	Book E
X	0 1 1 1 1 1		1 1 1 0 0 1 0 0 1 0 ?	<b>tlbsx</b>	TLB Search Indexed	Book E
X	0 1 1 1 1 1		1 1 1 0 0 1 0 1 1 0 /	<b>sthbrx</b>	Store Halfword Byte-Reverse Indexed	Book E
X	0 1 1 1 1 1		1 1 1 0 0 1 1 0 1 0 0	<b>extsh</b>	Extend Sign Halfword	Book E
X	0 1 1 1 1 1		1 1 1 0 0 1 1 0 1 0 1	<b>extsh.</b>	Extend Sign Halfword & record CR	Book E
X	0 1 1 1 1 1		1 1 1 0 1 1 0 0 1 0 /	<b>tlbre</b>	TLB Read Entry	Book E
X	0 1 1 1 1 1		1 1 1 0 1 1 1 0 1 0 0	<b>extsb</b>	Extend Sign Byte	Book E
X	0 1 1 1 1 1		1 1 1 0 1 1 1 0 1 0 1	<b>extsb.</b>	Extend Sign Byte & record CR	Book E
X	0 1 1 1 1 1		1 1 1 1 0 0 1 0 1 1 0	<b>divwuo</b>	Divide Word Unsigned & record OV	Book E
X	0 1 1 1 1 1		1 1 1 1 0 0 1 0 1 1 1	<b>divwuo.</b>	Divide Word Unsigned & record OV & CR	Book E
X	0 1 1 1 1 1		1 1 1 1 0 1 0 0 1 0 /	<b>tlbwe</b>	TLB Write Entry	Book E
X	0 1 1 1 1 1		1 1 1 1 0 1 0 1 1 0 /	<b>icbi</b>	Instruction Cache Block Invalidate Indexed	Book E
X	0 1 1 1 1 1		1 1 1 1 1 0 1 0 1 1 0	<b>divwo</b>	Divide Word & Record OV	Book E

Table 262. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		1 1 1 1 1 0 1 0 1 1 1	<b>divwo.</b>	Divide Word & Record OV & CR	Book E
X	0 1 1 1 1 1		1 1 1 1 1 1 0 1 1 1 0 /	<b>dcbz</b>	Data Cache Block Set to Zero Indexed	Book E
X	1 1 1 1 - -			Reserved		

## 15.2 Instruction index sorted by mnemonic

Table 263 lists all of the 16-bit VLE instructions, sorted by mnemonic.

**Table 263. 16-Bit VLE instructions sorted by mnemonic**

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
RR	0 0 0 0 0 1 0 0 y y y y x x x x	<b>se_add</b>	Add	<a href="#">on page 741</a>
IM5	0 0 1 0 0 0 0 i i i i i x x x x	<b>se_addi</b>	Add Immediate	<a href="#">on page 741</a>
RR	0 1 0 0 0 1 1 0 y y y y x x x x	<b>se_and</b>	AND	<a href="#">on page 743</a>
RR	0 1 0 0 0 1 1 1 y y y y x x x x	<b>se_and.</b>	AND and Record	<a href="#">on page 743</a>
RR	0 1 0 0 0 1 0 1 y y y y x x x x	<b>se_andc</b>	AND with Complement	<a href="#">on page 743</a>
IM5	0 0 1 0 1 1 1 i i i i i x x x x	<b>se_andi</b>	And Immediate	<a href="#">on page 743</a>
B8	1 1 1 0 0 o i i d d d d d d d d	<b>se_bc</b>	Branch Conditional	<a href="#">on page 745</a>
IM5	0 1 1 0 0 0 0 i i i i i x x x x	<b>se_bclri</b>	Bit Clear Immediate	<a href="#">on page 746</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 0	<b>se_bctr</b>	Branch to Count Register	<a href="#">on page 746</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1	<b>se_bctrl</b>	Branch to Count Register & Link	<a href="#">on page 746</a>
IM5	0 1 1 0 0 0 1 i i i i i x x x x	<b>se_bgeni</b>	Bit Generate Immediate	<a href="#">on page 746</a>
B8	1 1 1 0 1 0 0 1 d d d d d d d d	<b>se_bl</b>	Branch and Link	<a href="#">on page 748</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0	<b>se_blr</b>	Branch to Link Register	<a href="#">on page 747</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1	<b>se_blrl</b>	Branch to Link Register & Link	<a href="#">on page 747</a>

Table 263. 16-Bit VLE instructions sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
IM5	0 0 1 0 1 1 0 i i i i x x x x	<b>se_bmaski</b>	Bit Mask Generate Immediate	<a href="#">on page 748</a>
B8	1 1 1 0 1 0 0 0 d d d d d d d d	<b>se_b</b>	Branch	<a href="#">on page 744</a>
IM5	0 1 1 0 0 1 0 i i i i x x x x	<b>se_bseti</b>	Bit Set Immediate	<a href="#">on page 748</a>
IM5	0 1 1 0 0 1 1 i i i i x x x x	<b>se_btsti</b>	Bit Test Immediate	<a href="#">on page 748</a>
RR	0 0 0 0 1 1 0 0 y y y y x x x x	<b>se_cmp</b>	Compare	<a href="#">on page 749</a>
RR	0 0 0 0 1 1 1 0 y y y y x x x x	<b>se_cmph</b>	Compare Halfword	<a href="#">on page 750</a>
RR	0 0 0 0 1 1 1 1 y y y y x x x x	<b>se_cmphi</b>	Compare Halfword Logical	<a href="#">on page 751</a>
IM5	0 0 1 0 1 0 1 i i i i x x x x	<b>se_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
RR	0 0 0 0 1 1 0 1 y y y y x x x x	<b>se_cmpl</b>	Compare Logical	<a href="#">on page 752</a>
IM5	0 0 1 0 0 0 1 i i i i x x x x	<b>se_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
R	0 0 0 0 0 0 0 0 1 1 0 1 x x x x	<b>se_extsb</b>	Extend Sign Byte	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 0 1 1 1 1 x x x x	<b>se_extsh</b>	Extend Sign Halfword	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 0 1 1 0 0 x x x x	<b>se_extzb</b>	Extend with Zeros Byte	<a href="#">on page 756</a>
R	0 0 0 0 0 0 0 0 1 1 1 0 x x x x	<b>se_extzh</b>	Extend with Zeros Halfword	<a href="#">on page 756</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	<b>se_illegal</b>	Illegal	<a href="#">on page 757</a>

Table 263. 16-Bit VLE instructions sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1	<b>se_isync</b>	Instruction Synchronize	<a href="#">on page 757</a>
SD4	1 0 0 0 i i i i z z z z x x x x	<b>se_lbz</b>	Load Byte and Zero	<a href="#">on page 757</a>
SD4	1 0 1 0 i i i i z z z z x x x x	<b>se_lhz</b>	Load Halfword and Zero	<a href="#">on page 759</a>
IM7	0 1 0 0 1 i i i i i i i i x x x x	<b>se_li</b>	Load Immediate	<a href="#">on page 760</a>
SD4	1 1 0 0 i i i i z z z z x x x x	<b>se_lwz</b>	Load Word and Zero	<a href="#">on page 761</a>
RR	0 0 0 0 0 0 1 1 y y y y x x x x	<b>se_mfar</b>	Move from Alternate Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 1 0 x x x x	<b>se_mfctr</b>	Move From Count Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 0 0 x x x x	<b>se_mflr</b>	Move From Link Register	<a href="#">on page 762</a>
RR	0 0 0 0 0 0 0 1 y y y y x x x x	<b>se_mr</b>	Move Register	<a href="#">on page 762</a>
RR	0 0 0 0 0 0 1 0 y y y y x x x x	<b>se_mtar</b>	Move to Alternate Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 0 1 1 x x x x	<b>se_mtctr</b>	Move To Count Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 0 0 1 x x x x	<b>se_mtlr</b>	Move To Link Register	<a href="#">on page 763</a>
RR	0 0 0 0 0 1 0 1 y y y y x x x x	<b>se_mullw</b>	Multiply Low Word	<a href="#">on page 764</a>
R	0 0 0 0 0 0 0 0 0 0 1 1 x x x x	<b>se_neg</b>	Negate	<a href="#">on page 764</a>
R	0 0 0 0 0 0 0 0 0 0 1 0 x x x x	<b>se_not</b>	NOT	<a href="#">on page 764</a>



Table 263. 16-Bit VLE instructions sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
RR	0 1 0 0 0 1 0 0 y y y y x x x x	<b>se_or</b>	OR	<a href="#">on page 765</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 1	<b>se_rfc</b>	Return From Critical Interrupt	<a href="#">on page 767</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 0	<b>se_rfd</b>	Return From Debug Interrupt	<a href="#">on page 704</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0	<b>se_rfi</b>	Return From Interrupt	<a href="#">on page 767</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0	<b>se_sc</b>	System Call	<a href="#">on page 769</a>
RR	0 1 0 0 0 0 1 0 y y y y x x x x	<b>se_slw</b>	Shift Left Word	<a href="#">on page 770</a>
IM5	0 1 1 0 1 1 0 i i i i x x x x	<b>se_slwi</b>	Shift Left Word Immediate	<a href="#">on page 770</a>
RR	0 1 0 0 0 0 0 1 y y y y x x x x	<b>se_sraw</b>	Shift Right Algebraic Word	<a href="#">on page 771</a>
IM5	0 1 1 0 1 0 1 i i i i x x x x	<b>se_srawi</b>	Shift Right Algebraic Word Immediate	<a href="#">on page 771</a>
RR	0 1 0 0 0 0 0 0 y y y y x x x x	<b>se_srw</b>	Shift Right Word	<a href="#">on page 772</a>
IM5	0 1 1 0 1 0 0 i i i i x x x x	<b>se_srwi</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
SD4	1 0 0 1 i i i i z z z z x x x x	<b>se_stb</b>	Store Byte	<a href="#">on page 773</a>
SD4	1 0 1 1 i i i i z z z z x x x x	<b>se_sth</b>	Store Halfword	<a href="#">on page 773</a>
SD4	1 1 0 1 i i i i z z z z x x x x	<b>se_stw</b>	Store Word	<a href="#">on page 775</a>
RR	0 0 0 0 0 1 1 0 y y y y x x x x	<b>se_sub</b>	Subtract	<a href="#">on page 775</a>

**Table 263. 16-Bit VLE instructions sorted by mnemonic (continued)**

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
RR	0 0 0 0 0 1 1 1 y y y y x x x x	<b>se_subf</b>	Subtract From	<a href="#">on page 776</a>
IM5	0 0 1 0 0 1 0 i i i i i x x x x	<b>se_subi</b>	Subtract Immediate	<a href="#">on page 776</a>
IM5	0 0 1 0 0 1 1 i i i i i x x x x	<b>se_subi.</b>	Subtract Immediate and Record	<a href="#">on page 775</a>

[Table 264](#) outlines the 32-bit instruction encodings.

**Table 264. 32-bit instruction encodings (by mnemonic)**

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		01000010100	<b>add</b>	Add	Book E
X	011111		01000010101	<b>add.</b>	Add & record CR	Book E
D	000111	ttttt aaaaa iiii i	iiii iiii iiii i	<b>e_add16i</b>	Add Immediate	<a href="#">on page 741</a>
I16A	011100	iiii i aaaa 10001	iiii iiii iiii i	<b>e_add2i.</b>	Add (2 operand) Immediate and Record CR	<a href="#">on page 741</a>
I16A	011100	iiii i aaaa 10010	iiii iiii iiii i	<b>e_add2is</b>	Add (2 operand) Immediate Shifted	<a href="#">on page 741</a>
X	011111		00000010100	<b>addc</b>	Add Carrying	Book E
X	011111		00000010101	<b>addc.</b>	Add Carrying & record CR	Book E
X	011111		10000010100	<b>addco</b>	Add Carrying & record OV	Book E
X	011111		10000010101	<b>addco.</b>	Add Carrying & record OV & CR	Book E
X	011111		00100010100	<b>adde</b>	Add Extended with CA	Book E
X	011111		00100010101	<b>adde.</b>	Add Extended with CA & record CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		101000101000	<b>addeo</b>	Add Extended with CA & record OV	Book E
X	011111		101000101001	<b>addeo.</b>	Add Extended with CA & record OV & CR	Book E
SCI8	000110	ttttt aaaaa 10000	FSSiiiiiii	<b>e_addi</b>	Add Immediate	<a href="#">on page 741</a>
SCI8	000110	ttttt aaaaa 10001	FSSiiiiiii	<b>e_addi.</b>	Add Immediate and Record	<a href="#">on page 741</a>
SCI8	000110	ttttt aaaaa 10010	FSSiiiiiii	<b>e_addic</b>	Add Immediate Carrying	<a href="#">on page 743</a>
SCI8	000110	ttttt aaaaa 10011	FSSiiiiiii	<b>e_addic.</b>	Add Immediate Carrying and Record	<a href="#">on page 743</a>
X	011111		001110101000	<b>addme</b>	Add to Minus One Extended with CA	Book E
X	011111		001110101001	<b>addme.</b>	Add to Minus One Extended with CA & record CR	Book E
X	011111		101110101000	<b>addmeo</b>	Add to Minus One Extended with CA & record OV	Book E
X	011111		101110101001	<b>addmeo.</b>	Add to Minus One Extended with CA & record OV & CR	Book E
X	011111		110000101000	<b>addo</b>	Add & record OV	Book E
X	011111		110000101001	<b>addo.</b>	Add & record OV & CR	Book E
X	011111		001100101000	<b>addze</b>	Add to Zero Extended with CA	Book E
X	011111		001100101001	<b>addze.</b>	Add to Zero Extended with CA & record CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		10110010100	<b>addzeo</b>	Add to Zero Extended with CA & record OV	Book E
X	011111		10110010101	<b>addzeo.</b>	Add to Zero Extended with CA & record OV & CR	Book E
X	011111		00000111000	<b>and</b>	AND	Book E
X	011111		00000111001	<b>and.</b>	AND & record CR	Book E
I16L	011100	ttttt iiii 11001	iiii iiii i	<b>e_and2i.</b>	AND (2 operand) Immediate & record CR	<a href="#">on page 743</a>
I16L	011100	ttttt iiii 11101	iiii iiii i	<b>e_and2is.</b>	AND (2 operand) Immediate Shifted & record CR	<a href="#">on page 743</a>
X	011111		00001111000	<b>andc</b>	AND with Complement	Book E
X	011111		00001111001	<b>andc.</b>	AND with Complement & record CR	Book E
SCI8	000110	sssss aaaa 11000	FSSi i i i i i i	<b>e_andi</b>	AND Immediate	<a href="#">on page 743</a>
SCI8	000110	sssss aaaa 11001	FSSi i i i i i i	<b>e_andi.</b>	AND Immediate and Record	<a href="#">on page 743</a>
APU	00010-	- - - - -	- - - - -	<b>apu</b>	Reserved for APUs	
BD24	011110	0 d d d d d d d d d d	d d d d d d d d d d 0	<b>e_b</b>	Branch	<a href="#">on page 744</a>
BD15	011110	1000o o i i i i d d d d d	d d d d d d d d d d 0	<b>e_bc</b>	Branch Conditional	<a href="#">on page 745</a>
BD15	011110	1000o o i i i i d d d d d	d d d d d d d d d d 1	<b>e_bcl</b>	Branch Conditional & Link	<a href="#">on page 745</a>

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
BD24	011110	0 d d d d d d d d d d d d d d d d	d d d d d d d d d d d d 1	<b>e_bl</b>	Branch & Link	<a href="#">on page 744</a>
X	011111		0 0 0 0 0 0 0 0 0 0 0 /	<b>cmp</b>	Compare	Book E
I16A	011100	i i i i i a a a a a 1 0 0 1 1	i i i i i i i i i i i i	<b>e_cmp16i</b>	Compare Immediate	<a href="#">on page 749</a>
X	011111		0 0 0 0 0 0 1 1 1 1 0 /	<b>e_cmph</b>	Compare Halfword	<a href="#">on page 750</a>
I16A	011100	i i i i i a a a a a 1 0 1 1 0	i i i i i i i i i i i i	<b>e_cmph16i</b>	Compare Halfword Immediate	<a href="#">on page 750</a>
X	011111		0 0 0 0 1 0 1 1 1 1 0 /	<b>e_cmphl</b>	Compare Halfword Logical	<a href="#">on page 751</a>
I16A	011100	i i i i i a a a a a 1 0 1 1 1	i i i i i i i i i i i i	<b>e_cmphl16i</b>	Compare Halfword Logical Immediate	<a href="#">on page 751</a>
SCI8	000110	0 0 0 b f a a a a a 1 0 1 0 1	F S S i i i i i i i i	<b>e_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
X	011111		0 0 0 0 1 0 0 0 0 0 /	<b>cmpl</b>	Compare Logical	Book E
I16A	011100	i i i i i a a a a a 1 0 1 0 1	i i i i i i i i i i i i	<b>e_cmpl16i</b>	Compare Logical Immediate	<a href="#">on page 752</a>
SCI8	000110	0 0 1 b f a a a a a 1 0 1 0 1	F S S i i i i i i i i	<b>e_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
X	011111		0 0 0 0 0 1 1 0 1 0 0	<b>cntlzw</b>	Count Leading Zeros Word	Book E
X	011111		0 0 0 0 0 1 1 0 1 0 1	<b>cntlzw.</b>	Count Leading Zeros Word & record CR	Book E
XL	011111		0 1 0 0 0 0 0 0 0 1 /	<b>e_crand</b>	Condition Register AND	<a href="#">on page 753</a>
XL	011111		0 0 1 0 0 0 0 0 0 1 /	<b>e_crandc</b>	Condition Register AND with Complement	<a href="#">on page 753</a>



Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
XL	0 1 1 1 1 1		0 1 0 0 1 0 0 0 0 1 /	<b>e_creqv</b>	Condition Register Equivalent	<a href="#">on page 753</a>
XL	0 1 1 1 1 1		0 0 1 1 1 0 0 0 0 1 /	<b>e_crnand</b>	Condition Register NAND	Book E
XL	0 1 1 1 1 1		0 0 0 0 1 0 0 0 0 1 /	<b>e_crnor</b>	Condition Register NOR	<a href="#">on page 754</a>
XL	0 1 1 1 1 1		0 1 1 1 0 0 0 0 0 1 /	<b>e_cror</b>	Condition Register OR	<a href="#">on page 755</a>
XL	0 1 1 1 1 1		0 1 1 0 1 0 0 0 0 1 /	<b>e_crorc</b>	Condition Register OR with Complement	<a href="#">on page 755</a>
XL	0 1 1 1 1 1		0 0 1 1 0 0 0 0 0 1 /	<b>e_crxor</b>	Condition Register XOR	<a href="#">on page 755</a>
X	0 1 1 1 1 1		1 0 1 1 1 1 0 1 1 0 /	<b>dcba</b>	Data Cache Block Allocate Indexed	Book E
X	0 1 1 1 1 1		0 0 0 1 0 1 0 1 1 0 /	<b>dcbf</b>	Data Cache Block Flush Indexed	Book E
X	0 1 1 1 1 1		0 1 1 1 0 1 0 1 1 0 /	<b>dcbi</b>	Data Cache Block Invalidate Indexed	Book E
X	0 1 1 1 1 1		0 0 0 0 1 1 0 1 1 0 /	<b>dcbst</b>	Data Cache Block Store Indexed	Book E
X	0 1 1 1 1 1		0 1 0 0 0 1 0 1 1 0 /	<b>dcbt</b>	Data Cache Block Touch Indexed	Book E
X	0 1 1 1 1 1		0 0 1 1 1 1 0 1 1 0 /	<b>dcbstst</b>	Data Cache Block Touch for Store Indexed	Book E
X	0 1 1 1 1 1		1 1 1 1 1 1 0 1 1 0 /	<b>dcbz</b>	Data Cache Block Set to Zero Indexed	Book E
X	0 1 1 1 1 1		0 1 1 1 1 0 1 0 1 1 0	<b>divw</b>	Divide Word	Book E
X	0 1 1 1 1 1		0 1 1 1 1 0 1 0 1 1 1	<b>divw.</b>	Divide Word & record CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		11111010110	<b>divwo</b>	Divide Word & Record OV	Book E
X	011111		11111010111	<b>divwo.</b>	Divide Word & Record OV & CR	Book E
X	011111		01111010110	<b>divwu</b>	Divide Word Unsigned	Book E
X	011111		01111010111	<b>divwu.</b>	Divide Word Unsigned & record CR	Book E
X	011111		11111010110	<b>divwuo</b>	Divide Word Unsigned & record OV	Book E
X	011111		11111010111	<b>divwuo.</b>	Divide Word Unsigned & record OV & CR	Book E
X	011111		01000111000	<b>eqv</b>	Equivalent	Book E
X	011111		01000111001	<b>eqv.</b>	Equivalent & record CR	Book E
X	011111		11101110100	<b>extsb</b>	Extend Sign Byte	Book E
X	011111		11101110101	<b>extsb.</b>	Extend Sign Byte & record CR	Book E
X	011111		11100110100	<b>extsh</b>	Extend Sign Halfword	Book E
X	011111		11100110101	<b>extsh.</b>	Extend Sign Halfword & record CR	Book E
X	011111		1111010110 /	<b>icbi</b>	Instruction Cache Block Invalidate Indexed	Book E
X	011111		0000010110 /	<b>icbt</b>	Instruction Cache Block Touch Indexed	Book E
X	011111		- - - - 01111 /	<b>isel</b>	Integer Select	Book E
D	001100	t t t t t a a a a a d d d d d	d d d d d d d d d d d	<b>e_lbz</b>	Load Byte & Zero	<a href="#">on page 757</a>



Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
D8	000110	ttttt aaaaa 00000	000dd ddddd d	<b>e_lbzu</b>	Load Byte & Zero with Update	<a href="#">on page 757</a>
X	011111		00011 10111 /	<b>lbzux</b>	Load Byte & Zero with Update Indexed	Book E
X	011111		00010 10111 /	<b>lbzx</b>	Load Byte & Zero Indexed	Book E
D	001110	ttttt aaaaa ddddd	dddd d d d d d d	<b>e_lha</b>	Load Halfword Algebraic	<a href="#">on page 758</a>
D8	000110	ttttt aaaaa 00000	011dd ddddd d	<b>e_lhau</b>	Load Halfword Algebraic With Update	<a href="#">on page 758</a>
X	011111		01011 10111 /	<b>lhaux</b>	Load Halfword Algebraic with Update Indexed	Book E
X	011111		01010 10111 /	<b>lhax</b>	Load Halfword Algebraic Indexed	Book E
X	011111		11000 10110 /	<b>lhbrx</b>	Load Halfword Byte-Reverse Indexed	Book E
D	010110	ttttt aaaaa ddddd	dddd d d d d d d	<b>e_lhz</b>	Load Halfword & Zero	<a href="#">on page 759</a>
D8	000110	ttttt aaaaa 00000	001dd ddddd d	<b>e_lhzu</b>	Load Halfword & Zero with Update	<a href="#">on page 759</a>
X	011111		01001 10111 /	<b>lhzux</b>	Load Halfword & Zero with Update Indexed	Book E
X	011111		01000 10111 /	<b>lhzx</b>	Load Halfword & Zero Indexed	Book E
LI20	011100	ttttt iiii 0iiii	iiii i i i i i i	<b>e_li</b>	Load Immediate	<a href="#">on page 760</a>
I16L	011100	ttttt iiii 11100	iiii i i i i i i	<b>e_lis</b>	Load Immediate Shifted	<a href="#">on page 760</a>
D8	000110	ttttt aaaaa 00001	000dd ddddd d	<b>e_lmw</b>	Load Multiple Word	<a href="#">on page 761</a>





Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		0000010100 /	<b>lwarx</b>	Load Word & Reserve Indexed	Book E
X	011111		1000010110 /	<b>lwbrx</b>	Load Word Byte-Reverse Indexed	Book E
D	010100	ttttt aaaaa dddd	dddd dddd d	<b>e_lwz</b>	Load Word & Zero	<a href="#">on page 761</a>
D8	000110	ttttt aaaaa 00000	010dd dddd d	<b>e_lwzu</b>	Load Word & Zero with Update	<a href="#">on page 761</a>
X	011111		0000110111 /	<b>lwzux</b>	Load Word & Zero with Update Indexed	Book E
X	011111		0000010111 /	<b>lwzx</b>	Load Word & Zero Indexed	Book E
X	011111		1101010110 /	<b>mbar</b>	Memory Barrier	Book E
XL	011111		0000010000 /	<b>e_mcrf</b>	Move Condition Register Field	<a href="#">on page 761</a>
X	011111		1000000000 /	<b>mcrxr</b>	Move to Condition Register from XER	Book E
X	011111		0100010011 /	<b>mfapidi</b>	Move From APID Indirect	Book E
X	011111		0000010011 /	<b>mfcrr</b>	Move From Condition Register	Book E
XFX	011111		0101000011 /	<b>mfocr</b>	Move From Device Control Register	Book E
X	011111		0001010011 /	<b>mfmsr</b>	Move From Machine State Register	Book E
XFX	011111		0101010011 /	<b>mfpsr</b>	Move From Special Purpose Register	Book E
X	011111		1001010110 /	<b>msync</b>	Memory Synchronize	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
XFX	0111111		0010010000 /	<b>mtcrf</b>	Move To Condition Register Fields	Book E
XFX	0111111		0111000011 /	<b>mtdcr</b>	Move To Device Control Register	Book E
X	0111111		0010010010 /	<b>mtmsr</b>	Move To Machine State Register	Book E
XFX	0111111		0111010011 /	<b>mtspr</b>	Move To Special Purpose Register	Book E
X	0111111		/0010010110	<b>mulhw</b>	Multiply High Word	Book E
X	0111111		/0010010111	<b>mulhw.</b>	Multiply High Word & record CR	Book E
X	0111111		/0000010110	<b>mulhwu</b>	Multiply High Word Unsigned	Book E
X	0111111		/0000010111	<b>mulhwu.</b>	Multiply High Word Unsigned & Record	Book E
I16A	011100	iiiiiaaaaa10100	iiiiiiiiiii	<b>e_mull2i</b>	Multiply Low Word (2 operand) Immediate	<a href="#">on page 763</a>
SCI8	000110	tttttaaaaa10100	FSSiiiiiii	<b>e_mulli</b>	Multiply Low Immediate	<a href="#">on page 763</a>
X	0111111		00111010110	<b>mullw</b>	Multiply Low Word	Book E
X	0111111		00111010111	<b>mullw.</b>	Multiply Low Word & record CR	Book E
X	0111111		10111010110	<b>mullwo</b>	Multiply Low Word & record OV	Book E
X	0111111		10111010111	<b>mullwo.</b>	Multiply Low Word & record OV & CR	Book E
X	0111111		01110111000	<b>nand</b>	NAND	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		01110111001	<b>nand.</b>	NAND & record CR	Book E
X	011111		00011010000	<b>neg</b>	Negate	Book E
X	011111		00011010001	<b>neg.</b>	Negate & record CR	Book E
X	011111		10011010000	<b>nego</b>	Negate & record OV	Book E
X	011111		10011010001	<b>nego.</b>	Negate & record OV & record CR	Book E
X	011111		00011111000	<b>nor</b>	NOR	Book E
X	011111		00011111001	<b>nor.</b>	NOR & record CR	Book E
X	011111		01101111000	<b>or</b>	OR	Book E
X	011111		01101111001	<b>or.</b>	OR & record CR	Book E
I16L	011100	ttttt iiii 11000	iiii iiii i	<b>e_or2i</b>	OR (2 operand) Immediate	<a href="#">on page 765</a>
I16L	011100	ttttt iiii 11010	iiii iiii i	<b>e_or2is</b>	OR (2 operand) Immediate Shifted	<a href="#">on page 765</a>
X	011111		01100111000	<b>orc</b>	OR with Complement	Book E
X	011111		01100111001	<b>orc.</b>	OR with Complement & record CR	Book E
SCI8	000110	sssss aaaa 11010	FSSi iiii i	<b>e_ori</b>	OR Immediate	<a href="#">on page 768</a>
SCI8	000110	sssss aaaa 11011	FSSi iiii i	<b>e_ori.</b>	OR Immediate and Record	<a href="#">on page 768</a>
X	1111 - -			Reserved		
X	011111		01000110000	<b>e_rlw</b>	Rotate Left Word	<a href="#">on page 768</a>

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1		0 1 0 0 0 1 1 0 0 0 1	<b>e_rlwi.</b>	Rotate Left Word & record CR	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 1 1 1 0 0 0 0	<b>e_rlwi</b>	Rotate Left Word Immediate	<a href="#">on page 768</a>
X	0 1 1 1 1 1		0 1 0 0 1 1 1 0 0 0 1	<b>e_rlwi.</b>	Rotate Left Word Immediate & record CR	<a href="#">on page 768</a>
RLWI	0 1 1 1 0 1	s s s s s a a a a a h h h h h	b b b b b e e e e e 0	<b>e_rlwimi</b>	Rotate Left Word Immed then Mask Insert	<a href="#">on page 769</a>
RLWI	0 1 1 1 0 1	s s s s s a a a a a h h h h h	b b b b b e e e e e 1	<b>e_rlwinm</b>	Rotate Left Word Immed then AND with Mask	<a href="#">on page 770</a>
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 0 0 0	<b>slw</b>	Shift Left Word	Book E
X	0 1 1 1 1 1		0 0 0 0 0 1 1 0 0 0 1	<b>slw.</b>	Shift Left Word & record CR	Book E
X	0 1 1 1 1 1		0 0 0 0 1 1 1 0 0 0 0	<b>e_slwi</b>	Shift Left Word Immediate	<a href="#">on page 761</a>
X	0 1 1 1 1 1		0 0 0 0 1 1 1 0 0 0 1	<b>e_slwi.</b>	Shift Left Word Immediate & record CR	<a href="#">on page 761</a>
X	0 1 1 1 1 1		1 1 0 0 0 1 1 0 0 0 0	<b>sraw</b>	Shift Right Algebraic Word	Book E
X	0 1 1 1 1 1		1 1 0 0 0 1 1 0 0 0 1	<b>sraw.</b>	Shift Right Algebraic Word & record CR	Book E
X	0 1 1 1 1 1		1 1 0 0 1 1 1 0 0 0 0	<b>srawi</b>	Shift Right Algebraic Word Immediate	Book E
X	0 1 1 1 1 1		1 1 0 0 1 1 1 0 0 0 1	<b>srawi.</b>	Shift Right Algebraic Word Immediate & record CR	Book E
X	0 1 1 1 1 1		1 0 0 0 0 1 1 0 0 0 0	<b>srw</b>	Shift Right Word	Book E
X	0 1 1 1 1 1		1 0 0 0 0 1 1 0 0 0 1	<b>srw.</b>	Shift Right Word & record CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		10001 11000 0	<b>e_srwi</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
X	011111		10001 11000 1	<b>e_srwi.</b>	Shift Right Word Immediate & record CR	<a href="#">on page 772</a>
D	001101	ttttt aaaaa dddd	dddd d d d d d d	<b>e_stb</b>	Store Byte	<a href="#">on page 773</a>
D8	000110	ttttt aaaaa 00000	100 d d d d d d d	<b>e_stbu</b>	Store Byte with Update	<a href="#">on page 773</a>
X	011111		00111 10111 /	<b>stbux</b>	Store Byte with Update Indexed	Book E
X	011111		00110 10111 /	<b>stbx</b>	Store Byte Indexed	Book E
D	010111	ttttt aaaaa dddd	dddd d d d d d d	<b>e_sth</b>	Store Halfword	<a href="#">on page 773</a>
X	011111		11100 10110 /	<b>sthbrx</b>	Store Halfword Byte-Reverse Indexed	Book E
D8	000110	ttttt aaaaa 00000	101 d d d d d d d	<b>e_sthu</b>	Store Halfword with Update	<a href="#">on page 773</a>
X	011111		01101 10111 /	<b>sthux</b>	Store Halfword with Update Indexed	Book E
X	011111		01100 10111 /	<b>sthx</b>	Store Halfword Indexed	Book E
D8	000110	ttttt aaaaa 00001	001 d d d d d d d	<b>e_stmw</b>	Store Multiple Word	<a href="#">on page 774</a>
X	011111		10110 10101 /	<b>stswi</b>	Store String Word Immediate	Book E
X	011111		10100 10101 /	<b>stswx</b>	Store String Word Indexed	Book E
D	010101	ttttt aaaaa dddd	dddd d d d d d d	<b>e_stw</b>	Store Word	<a href="#">on page 775</a>

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		10100 10110 /	<b>stwbrx</b>	Store Word Byte-Reverse Indexed	Book E
X	011111		00100 10110 1	<b>stwcx.</b>	Store Word Conditional Indexed & record CR	Book E
D8	000110	t t t t t a a a a a 00000	110 d d d d d d d	<b>e_stwu</b>	Store Word with Update	<i>on page 775</i>
X	011111		00101 10111 /	<b>stwux</b>	Store Word with Update Indexed	Book E
X	011111		00100 10111 /	<b>stwx</b>	Store Word Indexed	Book E
X	011111		00001 01000 0	<b>subf</b>	Subtract From	Book E
X	011111		00001 01000 1	<b>subf.</b>	Subtract From & record CR	Book E
X	011111		00000 01000 0	<b>subfc</b>	Subtract From Carrying	Book E
X	011111		00000 01000 1	<b>subfc.</b>	Subtract From Carrying & record CR	Book E
X	011111		10000 01000 0	<b>subfco</b>	Subtract From Carrying & record OV	Book E
X	011111		10000 01000 1	<b>subfco.</b>	Subtract From Carrying & record OV & CR	Book E
X	011111		00100 01000 0	<b>subfe</b>	Subtract From Extended with CA	Book E
X	011111		00100 01000 1	<b>subfe.</b>	Subtract From Extended with CA & record CR	Book E
X	011111		10100 01000 0	<b>subfeo</b>	Subtract From Extended with CA & record OV	Book E
X	011111		10100 01000 1	<b>subfeo.</b>	Subtract From Extended with CA & record OV & CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
SCI8	000110	ttttt aaaaa 10110	FSSiiiiiii	<b>e_subfic</b>	Subtract from Immediate Carrying	<a href="#">on page 776</a>
SCI8	000110	ttttt aaaaa 10111	FSSiiiiiii	<b>e_subfic.</b>	Subtract from Immediate and Record	<a href="#">on page 776</a>
X	011111		00111 01000 0	<b>subfme</b>	Subtract From Minus One Extended with CA	Book E
X	011111		00111 01000 1	<b>subfme.</b>	Subtract From Minus One Extended with CA & record CR	Book E
X	011111		10111 01000 0	<b>subfmeo</b>	Subtract From Minus One Extended with CA & record OV	Book E
X	011111		10111 01000 1	<b>subfmeo.</b>	Subtract From Minus One Extended with CA & record OV & CR	Book E
X	011111		10001 01000 0	<b>subfo</b>	Subtract From & record OV	Book E
X	011111		10001 01000 1	<b>subfo.</b>	Subtract From & record OV & CR	Book E
X	011111		00110 01000 0	<b>subfze</b>	Subtract From Zero Extended with CA	Book E
X	011111		00110 01000 1	<b>subfze.</b>	Subtract From Zero Extended with CA & record CR	Book E
X	011111		10110 01000 0	<b>subfzeo</b>	Subtract From Zero Extended with CA & record OV	Book E
X	011111		10110 01000 1	<b>subfzeo.</b>	Subtract From Zero Extended with CA & record OV & CR	Book E

Table 264. 32-bit instruction encodings (by mnemonic) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111		1100010010 /	<b>tlbivax</b>	TLB Invalidate Virtual Address Indexed	Book E
X	011111		1110110010 /	<b>tlbre</b>	TLB Read Entry	Book E
X	011111		1110010010 ?	<b>tlbsx</b>	TLB Search Indexed	Book E
X	011111		1000110110 /	<b>tlbsync</b>	TLB Synchronize	Book E
X	011111		1111010010 /	<b>tlbwe</b>	TLB Write Entry	Book E
X	011111		0000000100 /	<b>tw</b>	Trap Word	Book E
X	011111		0010000011 /	<b>wrtee</b>	Write External Enable	Book E
X	011111		0010100011 /	<b>wrteei</b>	Write External Enable Immediate	Book E
X	011111		0100111100 0	<b>xor</b>	XOR	Book E
X	011111		0100111100 1	<b>xor.</b>	XOR & record CR	Book E
SCI8	000110	sssss aaaaa 11100	FSSiiiiiii	<b>e_xori</b>	XOR Immediate	<a href="#">on page 777</a>
SCI8	000110	sssss aaaaa 11101	FSSiiiiiii	<b>e_xori.</b>	XOR Immediate and Record	<a href="#">on page 777</a>



### 15.3 Instruction index sorted by opcode

Table 265 lists all the 16-bit Power\*Embedded instructions, sorted by opcode.

Table 265. Instruction index sorted by opcode

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1	se_isync	Instruction Synchronize	<a href="#">on page 757</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0	se_sc	System Call	<a href="#">on page 769</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0	se_blr	Branch to Link Register	<a href="#">on page 747</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 1	se_blrI	Branch to Link Register & Link	<a href="#">on page 747</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 0	se_bctr	Branch to Count Register	<a href="#">on page 746</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1	se_bctrl	Branch to Count Register & Link	<a href="#">on page 746</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0	se_rfi	Return From Interrupt	<a href="#">on page 767</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 1	se_rfci	Return From Critical Interrupt	<a href="#">on page 767</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 0	se_rfdi	Return From Debug Interrupt	<a href="#">on page 704</a>
C	0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 1	unimp		
C	0 0 0 0 0 0 0 0 0 0 0 0 1 1 - -	unimp		
C	0 0 0 0 0 0 0 0 0 0 0 0 1 - - - -	unimp		
R	0 0 0 0 0 0 0 0 0 0 1 0 x x x x	se_not	NOT	<a href="#">on page 764</a>
R	0 0 0 0 0 0 0 0 0 0 1 1 x x x x	unimp		
R	0 0 0 0 0 0 0 0 0 1 0 0 x x x x	se_lmW	Load Multiple Word	<a href="#">on page 760</a>
R	0 0 0 0 0 0 0 0 0 1 0 1 x x x x	se_stmW	Store Multiple Word	<a href="#">on page 774</a>

Table 265. Instruction index sorted by opcode (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
R	0 0 0 0 0 0 0 0 0 1 1 - x x x x	unimp		
R	0 0 0 0 0 0 0 0 1 0 0 0 x x x x	se_mflr	Move From Link Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 0 1 x x x x	se_mtlr	Move To Link Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 0 1 0 x x x x	se_mfctr	Move From Count Register	<a href="#">on page 762</a>
R	0 0 0 0 0 0 0 0 1 0 1 1 x x x x	se_mtctr	Move To Count Register	<a href="#">on page 763</a>
R	0 0 0 0 0 0 0 0 1 1 0 0 x x x x	se_extzb	Extend with Zeros Byte	<a href="#">on page 756</a>
R	0 0 0 0 0 0 0 0 1 1 0 1 x x x x	se_extsb	Extend Sign Byte	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 0 1 1 1 0 x x x x	se_extzh	Extend with Zeros Halfword	<a href="#">on page 756</a>
R	0 0 0 0 0 0 0 0 1 1 1 1 x x x x	se_extsh	Extend Sign Halfword	<a href="#">on page 755</a>
R	0 0 0 0 0 0 0 1 - - - - x x x x	unimp		
RR	0 0 0 0 0 0 1 0 y y y y x x x x	se_mtar	Move to Alternate Register	<a href="#">on page 763</a>
RR	0 0 0 0 0 0 1 1 y y y y x x x x	se_mfar	Move from Alternate Register	<a href="#">on page 762</a>
RR	0 0 0 0 0 1 0 0 y y y y x x x x	se_add	Add	<a href="#">on page 741</a>
RR	0 0 0 0 0 1 0 1 y y y y x x x x	unimp		
RR	0 0 0 0 0 1 1 0 y y y y x x x x	se_sub	Subtract	<a href="#">on page 775</a>
RR	0 0 0 0 0 1 1 1 y y y y x x x x	se_sub.	Subtract and Record	<a href="#">on page 775</a>
RR	0 0 0 0 1 0 0 0 y y y y x x x x	unimp		
RR	0 0 0 0 1 0 0 1 y y y y x x x x	unimp		

Table 265. Instruction index sorted by opcode (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
RR	0 0 0 0 1 0 1 0 y y y y x x x x	<b>se_mullw</b>	Multiply Low Word	<a href="#">on page 764</a>
RR	0 0 0 0 1 0 1 1 y y y y x x x x	<b>se_mr</b>	Move Register	<a href="#">on page 762</a>
RR	0 0 0 0 1 1 0 0 y y y y x x x x	<b>se_cmp</b>	Compare	<a href="#">on page 749</a>
RR	0 0 0 0 1 1 0 1 y y y y x x x x	<b>se_cmpl</b>	Compare Logical	<a href="#">on page 752</a>
RR	0 0 0 0 1 1 1 0 y y y y x x x x	<b>se_xor</b>	XOR	<a href="#">on page 777</a>
RR	0 0 0 0 1 1 1 1 y y y y x x x x	<b>se_or</b>	OR	<a href="#">on page 765</a>
IM5	0 0 1 0 0 0 0 i i i i i x x x x	<b>se_addi</b>	Add Immediate	<a href="#">on page 741</a>
IM5	0 0 1 0 0 0 1 i i i i i x x x x	<b>se_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
IM5	0 0 1 0 0 1 0 i i i i i x x x x	<b>se_subi</b>	Subtract Immediate	<a href="#">on page 776</a>
IM5	0 0 1 0 0 1 1 i i i i i x x x x	<b>se_subi.</b>	Subtract Immediate and Record	<a href="#">on page 776</a>
IM5	0 0 1 0 1 0 0 i i i i i x x x x	<b>se_subfic</b>	Subtract From Immediate Carrying	<a href="#">on page 776</a>
IM5	0 0 1 0 1 0 1 i i i i i x x x x	<b>se_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
IM5	0 0 1 0 1 1 0 i i i i i x x x x	<b>se_bmask i</b>	Bit Mask Generate Immediate	<a href="#">on page 748</a>
IM5	0 0 1 0 1 1 1 i i i i i x x x x	<b>se_andi</b>	And Immediate	<a href="#">on page 743</a>
RR	0 1 0 0 0 0 0 0 y y y y x x x x	<b>se_sraw</b>	Shift Right Algebraic Word	<a href="#">on page 771</a>

Table 265. Instruction index sorted by opcode (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
RR	0 1 0 0 0 0 0 1 y y y y x x x x	<b>se_rlw</b>	Rotate Left Word	<a href="#">on page 768</a>
RR	0 1 0 0 0 0 1 0 y y y y x x x x	<b>se_srw</b>	Shift Right Word	<a href="#">on page 772</a>
RR	0 1 0 0 0 0 1 1 y y y y x x x x	<b>se_slw</b>	Shift Left Word	<a href="#">on page 770</a>
RR	0 1 0 0 0 1 0 0 y y y y x x x x	<b>se_subf</b>	Subtract From	<a href="#">on page 776</a>
RR	0 1 0 0 0 1 0 1 y y y y x x x x	<b>se_andc</b>	AND with Complement	<a href="#">on page 743</a>
RR	0 1 0 0 0 1 1 0 y y y y x x x x	<b>se_and</b>	AND	<a href="#">on page 743</a>
RR	0 1 0 0 0 1 1 1 y y y y x x x x	<b>se_and.</b>	AND and Record	<a href="#">on page 743</a>
IM7	0 1 0 0 1 i i i i i i i i x x x x	<b>se_li</b>	Load Immediate	<a href="#">on page 760</a>
IM5	0 1 1 0 0 0 0 i i i i i x x x x	<b>se_bclri</b>	Bit Clear Immediate	<a href="#">on page 746</a>
IM5	0 1 1 0 0 0 1 i i i i i x x x x	<b>se_bgeni</b>	Bit Generate Immediate	<a href="#">on page 747</a>
IM5	0 1 1 0 0 1 0 i i i i i x x x x	<b>se_bseti</b>	Bit Set Immediate	<a href="#">on page 748</a>
IM5	0 1 1 0 0 1 1 i i i i i x x x x	<b>se_btsti</b>	Bit Test Immediate	<a href="#">on page 748</a>
IM5	0 1 1 0 1 0 0 i i i i i x x x x	<b>se_rlwi</b>	Rotate Left Word Immediate	<a href="#">on page 768</a>
IM5	0 1 1 0 1 0 1 i i i i i x x x x	<b>se_srawi</b>	Shift Right Algebraic Word Immediate	<a href="#">on page 771</a>
IM5	0 1 1 0 1 1 0 i i i i i x x x x	<b>se_slwi</b>	Shift Left Word Immediate	<a href="#">on page 770</a>

Table 265. Instruction index sorted by opcode (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )	Mnemonic	Instruction	Page
IM5	0 1 1 0 1 1 1 i i i i x x x x	<b>se_srwi</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
SD4	1 0 0 0 i i i i z z z z x x x x	<b>se_lbz</b>	Load Byte and Zero	<a href="#">on page 757</a>
SD4	1 0 0 1 i i i i z z z z x x x x	<b>se_stb</b>	Store Byte	<a href="#">on page 773</a>
SD4	1 0 1 0 i i i i z z z z x x x x	<b>se_lhz</b>	Load Halfword and Zero	<a href="#">on page 759</a>
SD4	1 0 1 1 i i i i z z z z x x x x	<b>se_sth</b>	Store Halfword	<a href="#">on page 773</a>
SD4	1 1 0 0 i i i i z z z z x x x x	<b>se_lwz</b>	Load Word and Zero	<a href="#">on page 761</a>
SD4	1 1 0 1 i i i i z z z z x x x x	<b>se_stw</b>	Store Word	<a href="#">on page 775</a>
UNIMP	1 1 1 0 - - - - - - - -	unimp		
B8	1 1 1 1 0 o i i d d d d d d d d	<b>se_bc</b>	Branch Conditional	<a href="#">on page 745</a>
B8	1 1 1 1 1 0 0 0 d d d d d d d d	<b>se_b</b>	Branch	<a href="#">on page 744</a>
B8	1 1 1 1 1 0 0 1 d d d d d d d d	<b>se_bl</b>	Branch and Link	<a href="#">on page 744</a>
UNIMP	1 1 1 1 1 0 1 - - - - - - - -	unimp		
UNIMP	1 1 1 1 1 1 - - - - - - - -	unimp		

Table 266 outlines the 32-bit instruction encodings.

Table 266. 32-bit instruction encodings

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
APU	0-01--	-----	-----	<b>apu</b>	Reserved for APUs	
D14	001100	ttttt aaaaa 00ddd	dddd dddd d	<b>e_lbz</b>	Load Byte & Zero	<a href="#">on page 757</a>
D14	001100	ttttt aaaaa 01ddd	dddd dddd d	<b>e_lhz</b>	Load Halfword & Zero	<a href="#">on page 759</a>
D14	001100	ttttt aaaaa 10ddd	dddd dddd d	<b>e_lwz</b>	Load Word & Zero	<a href="#">on page 761</a>
D14	001100	ttttt aaaaa 11ddd	dddd dddd d	<b>e_ld</b>	Load Doubleword & Zero (reserved for 64b GPR)	
D14	001101	ttttt aaaaa 00ddd	dddd dddd d	<b>e_stb</b>	Store Byte	<a href="#">on page 773</a>
D14	001101	ttttt aaaaa 01ddd	dddd dddd d	<b>e_sth</b>	Store Halfword	<a href="#">on page 773</a>
D14	001101	ttttt aaaaa 10ddd	dddd dddd d	<b>e_stw</b>	Store Word	<a href="#">on page 775</a>
D14	001101	ttttt aaaaa 11ddd	dddd dddd d	<b>e_std</b>	Store Doubleword (reserved for 64b GPR)	
D8	001110	ttttt aaaaa 00000	000dd dddd d	<b>e_lbzu</b>	Load Byte & Zero with Update	<a href="#">on page 757</a>
D8	001110	ttttt aaaaa 00000	001dd dddd d	<b>e_lhzu</b>	Load Halfword & Zero with Update	<a href="#">on page 759</a>
D8	001110	ttttt aaaaa 00000	010dd dddd d	<b>e_lwzu</b>	Load Word & Zero with Update	<a href="#">on page 761</a>
D8	001110	ttttt aaaaa 00000	011dd dddd d	<b>e_ldu</b>	Load Doubleword with Update (reserved for 64b GPR)	
D8	001110	ttttt aaaaa 00000	100dd dddd d	<b>e_stbu</b>	Store Byte with Update	<a href="#">on page 773</a>



Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
D8	001110	ttttt aaaaa 00000	101dd ddddd d	e_sthu	Store Halfword with Update	<a href="#">on page 773</a>
D8	001110	ttttt aaaaa 00000	110dd ddddd d	e_stwu	Store Word with Update	<a href="#">on page 775</a>
D8	001110	ttttt aaaaa 00000	111dd ddddd d	e_stdu	Store Doubleword with Update (reserved for 64b GPR)	
D8	001110	ttttt aaaaa 00001	000dd ddddd d	e_lmww	Load Multiple Word	<a href="#">on page 760</a>
D8	001110	ttttt aaaaa 00001	001dd ddddd d	e_stmw	Store Multiple Word	<a href="#">on page 774</a>
D8	001110	ttttt aaaaa 00001	010dd ddddd d	e_lha	Load Halfword Algebraic	<a href="#">on page 758</a>
D8	001110	ttttt aaaaa 00001	011dd ddddd d	e_lhau	Load Halfword Algebraic with Update	<a href="#">on page 758</a>
UNIMP	001110	ttttt aaaaa 00001	1- - dd ddddd d	unimp		
UNIMP	001110	ttttt aaaaa 0001	-----	unimp		
UNIMP	001110	ttttt aaaaa 001	-----	unimp		
SCI8	001110	00000 aaaaa 010bF	Fssii iiii i	e_cmpi	Compare Immediate	<a href="#">on page 749</a>
SCI8	001110	00000 aaaaa 011bF	Fssii iiii i	e_cmpli	Compare Logical Immediate	<a href="#">on page 752</a>
SCI8	001110	ttttt aaaaa 10000	Fssii iiii i	e_addi	Add Immediate	<a href="#">on page 741</a>
SCI8	001110	ttttt aaaaa 10001	Fssii iiii i	e_addic	Add Immediate Carrying	<a href="#">on page 743</a>
SCI8	001110	ttttt aaaaa 10010	Fssii iiii i	e_andi	AND Immediate	<a href="#">on page 743</a>
SCI8	001110	ttttt aaaaa 10011	Fssii iiii i	e_ori	OR Immediate	<a href="#">on page 768</a>



Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
SCI8	001110	tttttt aaaaa 10100	Fssii i i i i i i	<b>e_subfic</b>	Subtract from Immediate Carrying	<a href="#">on page 776</a>
SCI8	001110	tttttt aaaaa 10101	Fssii i i i i i i	unimp		
SCI8	001110	tttttt aaaaa 10110	Fssii i i i i i i	<b>e_mulli</b>	Multiply Low Immediate	<a href="#">on page 763</a>
SCI8	001110	tttttt aaaaa 10111	Fssii i i i i i i	<b>e_xori</b>	XOR Immediate	<a href="#">on page 777</a>
SCI8	001110	tttttt aaaaa 11000	Fssii i i i i i i	<b>e_addi.</b>	Add Immediate and Record	<a href="#">on page 741</a>
SCI8	001110	tttttt aaaaa 11001	Fssii i i i i i i	<b>e_addic.</b>	Add Immediate Carrying and Record	<a href="#">on page 743</a>
SCI8	001110	tttttt aaaaa 11010	Fssii i i i i i i	<b>e_andi.</b>	AND Immediate and Record	<a href="#">on page 743</a>
SCI8	001110	tttttt aaaaa 11011	Fssii i i i i i i	<b>e_ori.</b>	OR Immediate and Record	<a href="#">on page 768</a>
SCI8	001110	tttttt aaaaa 11100	Fssii i i i i i i	<b>e_subfic.</b>	Subtract from Immediate and Record	<a href="#">on page 776</a>
SCI8	001110	tttttt aaaaa 10101	Fssii i i i i i i	unimp		
SCI8	001110	tttttt aaaaa 11110	Fssii i i i i i i	<b>e_mulli.</b>	Multiply Low Immediate and Record	<a href="#">on page 763</a>
SCI8	001110	tttttt aaaaa 11111	Fssii i i i i i i	<b>e_xori.</b>	XOR Immediate and Record	<a href="#">on page 777</a>
D	001111	tttttt aaaaa i i i i i	i i i i i i i i i i i	<b>e_add16i</b>	Add Immediate	<a href="#">on page 741</a>
LI20	011100	tttttt 0 i i i i i i i i i	i i i i i i i i i i i	<b>e_li</b>	Load Immediate	<a href="#">on page 760</a>
LI20	011100	tttttt 1 i i i i i i i i i	i i i i i i i i i i i	<b>e_lis</b>	Load Immediate Shifted	<a href="#">on page 760</a>





Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
RLWI	0111101	sssss aaaaa hhhh	bbbb eeeee 0	<b>e_rlwimi.</b>	Rotate Left Word Immed then Mask Insert & record CR	<a href="#">on page 768</a>
RLWI	0111101	sssss aaaaa hhhh	bbbb eeeee 1	<b>e_rlwinm.</b>	Rotate Left Word Immed then AND with Mask & record CR	<a href="#">on page 769</a>
BD24	0111110	0dddd dddd dddd	dddd dddd 0	<b>e_b</b>	Branch	<a href="#">on page 744</a>
BD24	0111110	0dddd dddd dddd	dddd dddd 1	<b>e_bl</b>	Branch & Link	<a href="#">on page 744</a>
BD15	0111110	1000o oiiii dddd	dddd dddd 0	<b>e_bc</b>	Branch Conditional	<a href="#">on page 745</a>
BD15	0111110	1000o oiiii dddd	dddd dddd 1	<b>e_bcl</b>	Branch Conditional & Link	<a href="#">on page 745</a>
X	0111111	-----	----- 01111 /	<b>isel</b>	Integer Select	Book E
X	0111111	-----	/ 0000 01011 0	<b>mulhwu</b>	Multiply High Word Unsigned	Book E
X	0111111	-----	/ 0000 01011 1	<b>mulhwu.</b>	Multiply High Word Unsigned & Record	Book E
X	0111111	-----	/ 0010 01011 0	<b>mulhw</b>	Multiply High Word	Book E
X	0111111	-----	/ 0010 01011 1	<b>mulhw.</b>	Multiply High Word & record CR	Book E
X	0111111	-----	00000 00000 /	<b>cmp</b>	Compare	Book E
X	0111111	-----	00000 00100 /	<b>tw</b>	Trap Word	Book E
X	0111111	-----	00000 01000 0	<b>subfc</b>	Subtract From Carrying	Book E
X	0111111	-----	00000 01000 1	<b>subfc.</b>	Subtract From Carrying & record CR	Book E
X	0111111	-----	00000 01010 0	<b>addc</b>	Add Carrying	Book E
X	0111111	-----	00000 01010 1	<b>addc.</b>	Add Carrying & record CR	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	00000 10011 /	<b>mfcrr</b>	Move From Condition Register	Book E
X	011111	-----	00000 10100 /	<b>lwarx</b>	Load Word & Reserve Indexed	Book E
X	011111	-----	00000 10110 /	<b>icbt</b>	Instruction Cache Block Touch Indexed	Book E
X	011111	-----	00000 10111 /	<b>lwzx</b>	Load Word & Zero Indexed	Book E
X	011111	-----	00000 11000 0	<b>slw</b>	Shift Left Word	Book E
X	011111	-----	00000 11000 1	<b>slw.</b>	Shift Left Word & record CR	Book E
X	011111	-----	00000 11010 0	<b>cntlzw</b>	Count Leading Zeros Word	Book E
X	011111	-----	00000 11010 1	<b>cntlzw.</b>	Count Leading Zeros Word & record CR	Book E
X	011111	-----	00000 11100 0	<b>and</b>	AND	Book E
X	011111	-----	00000 11100 1	<b>and.</b>	AND & record CR	Book E
X	011111	-----	00001 00000 /	<b>cmpl</b>	Compare Logical	Book E
X	011111	-----	00001 01000 0	<b>subf</b>	Subtract From	Book E
X	011111	-----	00001 01000 1	<b>subf.</b>	Subtract From & record CR	Book E
X	011111	-----	00001 10110 /	<b>dcbst</b>	Data Cache Block Store Indexed	Book E
X	011111	-----	00001 10111 /	<b>lwzux</b>	Load Word & Zero with Update Indexed	Book E
X	011111	-----	00001 11100 0	<b>andc</b>	AND with Complement	Book E
X	011111	-----	00001 11100 1	<b>andc.</b>	AND with Complement & record CR	Book E
X	011111	-----	00010 10011 /	<b>mfmsr</b>	Move From Machine State Register	Book E
X	011111	-----	00010 10110 /	<b>dcbf</b>	Data Cache Block Flush Indexed	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	00010 10111 /	<b>lbzx</b>	Load Byte & Zero Indexed	Book E
X	011111	-----	00011 01000 0	<b>neg</b>	Negate	Book E
X	011111	-----	00011 01000 1	<b>neg.</b>	Negate & record CR	Book E
X	011111	-----	00011 10111 /	<b>lbzux</b>	Load Byte & Zero with Update Indexed	Book E
X	011111	-----	00011 11100 0	<b>nor</b>	NOR	Book E
X	011111	-----	00011 11100 1	<b>nor.</b>	NOR & record CR	Book E
X	011111	-----	00100 00011 /	<b>wrtee</b>	Write External Enable	Book E
X	011111	-----	00100 01000 0	<b>subfe</b>	Subtract From Extended with CA	Book E
X	011111	-----	00100 01000 1	<b>subfe.</b>	Subtract From Extended with CA & record CR	Book E
X	011111	-----	00100 01010 0	<b>adde</b>	Add Extended with CA	Book E
X	011111	-----	00100 01010 1	<b>adde.</b>	Add Extended with CA & record CR	Book E
XFX	011111	-----	00100 10000 /	<b>mtrcf</b>	Move To Condition Register Fields	Book E
X	011111	-----	00100 10010 /	<b>mtmsr</b>	Move To Machine State Register	Book E
X	011111	-----	00100 10110 1	<b>stwcx.</b>	Store Word Conditional Indexed & record CR	Book E
X	011111	-----	00100 10111 /	<b>stwx</b>	Store Word Indexed	Book E
X	011111	-----	00101 00011 /	<b>wrteei</b>	Write External Enable Immediate	Book E
X	011111	-----	00101 10111 /	<b>stwux</b>	Store Word with Update Indexed	Book E
X	011111	-----	00110 01000 0	<b>subfze</b>	Subtract From Zero Extended with CA	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	00110 01000 1	<b>subfze.</b>	Subtract From Zero Extended with CA & record CR	Book E
X	011111	-----	00110 01010 0	<b>addze</b>	Add to Zero Extended with CA	Book E
X	011111	-----	00110 01010 1	<b>addze.</b>	Add to Zero Extended with CA & record CR	Book E
X	011111	-----	00110 10111 /	<b>stbx</b>	Store Byte Indexed	Book E
X	011111	-----	00111 01000 0	<b>subfme</b>	Subtract From Minus One Extended with CA	Book E
X	011111	-----	00111 01000 1	<b>subfme.</b>	Subtract From Minus One Extended with CA & record CR	Book E
X	011111	-----	00111 01010 0	<b>addme</b>	Add to Minus One Extended with CA	Book E
X	011111	-----	00111 01010 1	<b>addme.</b>	Add to Minus One Extended with CA & record CR	Book E
X	011111	-----	00111 01011 0	<b>mullw</b>	Multiply Low Word	Book E
X	011111	-----	00111 01011 1	<b>mullw.</b>	Multiply Low Word & record CR	Book E
X	011111	-----	00111 10110 /	<b>dcbtst</b>	Data Cache Block Touch for Store Indexed	Book E
X	011111	-----	00111 10111 /	<b>stbux</b>	Store Byte with Update Indexed	Book E
X	011111	-----	01000 01010 0	<b>add</b>	Add	Book E
X	011111	-----	01000 01010 1	<b>add.</b>	Add & record CR	Book E
X	011111	-----	01000 10011 /	<b>mfapidi</b>	Move From APID Indirect	Book E
X	011111	-----	01000 10110 /	<b>dcbt</b>	Data Cache Block Touch Indexed	Book E
X	011111	-----	01000 10111 /	<b>lhzx</b>	Load Halfword & Zero Indexed	Book E
X	011111	-----	01000 11100 0	<b>eqv</b>	Equivalent	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	01000 11100 1	<b>eqv.</b>	Equivalent & record CR	Book E
X	011111	-----	01001 10111 /	<b>lhzux</b>	Load Halfword & Zero with Update Indexed	Book E
X	011111	-----	01001 11100 0	<b>xor</b>	XOR	Book E
X	011111	-----	01001 11100 1	<b>xor.</b>	XOR & record CR	Book E
XFX	011111	-----	01010 00011 /	<b>mfdcr</b>	Move From Device Control Register	Book E
XFX	011111	-----	01010 10011 /	<b>mfspr</b>	Move From Special Purpose Register	Book E
X	011111	-----	01010 10111 /	<b>lhax</b>	Load Halfword Algebraic Indexed	Book E
X	011111	-----	01011 10111 /	<b>lhaux</b>	Load Halfword Algebraic with Update Indexed	Book E
X	011111	-----	01100 10111 /	<b>sthx</b>	Store Halfword Indexed	Book E
X	011111	-----	01100 11100 0	<b>orc</b>	OR with Complement	Book E
X	011111	-----	01100 11100 1	<b>orc.</b>	OR with Complement & record CR	Book E
X	011111	-----	01101 10111 /	<b>sthux</b>	Store Halfword with Update Indexed	Book E
X	011111	-----	01101 11100 0	<b>or</b>	OR	Book E
X	011111	-----	01101 11100 1	<b>or.</b>	OR & record CR	Book E
XFX	011111	-----	01110 00011 /	<b>mtdcr</b>	Move To Device Control Register	Book E
X	011111	-----	01110 01011 0	<b>divwu</b>	Divide Word Unsigned	Book E
X	011111	-----	01110 01011 1	<b>divwu.</b>	Divide Word Unsigned & record CR	Book E
XFX	011111	-----	01110 10011 /	<b>mtspr</b>	Move To Special Purpose Register	Book E
X	011111	-----	01110 10110 /	<b>dcbi</b>	Data Cache Block Invalidate Indexed	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	01110 11100 0	<b>nand</b>	NAND	Book E
X	011111	-----	01110 11100 1	<b>nand.</b>	NAND & record CR	Book E
X	011111	-----	01111 01011 0	<b>divw</b>	Divide Word	Book E
X	011111	-----	01111 01011 1	<b>divw.</b>	Divide Word & record CR	Book E
X	011111	-----	10000 00000 /	<b>mcrxr</b>	Move to Condition Register from XER	Book E
X	011111	-----	10000 01000 0	<b>subfco</b>	Subtract From Carrying & record OV	Book E
X	011111	-----	10000 01000 1	<b>subfco.</b>	Subtract From Carrying & record OV & CR	Book E
X	011111	-----	10000 01010 0	<b>addco</b>	Add Carrying & record OV	Book E
X	011111	-----	10000 01010 1	<b>addco.</b>	Add Carrying & record OV & CR	Book E
X	011111	-----	10000 10110 /	<b>lwbrx</b>	Load Word Byte-Reverse Indexed	Book E
X	011111	-----	10000 11000 0	<b>srw</b>	Shift Right Word	Book E
X	011111	-----	10000 11000 1	<b>srw.</b>	Shift Right Word & record CR	Book E
X	011111	-----	10001 01000 0	<b>subfo</b>	Subtract From & record OV	Book E
X	011111	-----	10001 01000 1	<b>subfo.</b>	Subtract From & record OV & CR	Book E
X	011111	-----	10001 10110 /	<b>tlbsync</b>	TLB Synchronize	Book E
X	011111	-----	10010 10110 /	<b>msync</b>	Memory Synchronize	Book E
X	011111	-----	10011 01000 0	<b>nego</b>	Negate & record OV	Book E
X	011111	-----	10011 01000 1	<b>nego.</b>	Negate & record OV & record CR	Book E
X	011111	-----	10100 01000 0	<b>subfeo</b>	Subtract From Extended with CA & record OV	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	10100 01000 1	<b>subfeo.</b>	Subtract From Extended with CA & record OV & CR	Book E
X	011111	-----	10100 01010 0	<b>addeo</b>	Add Extended with CA & record OV	Book E
X	011111	-----	10100 01010 1	<b>addeo.</b>	Add Extended with CA & record OV & CR	Book E
X	011111	-----	10100 10101 /	<b>stswx</b>	Store String Word Indexed	Book E
X	011111	-----	10100 10110 /	<b>stwbrx</b>	Store Word Byte-Reverse Indexed	Book E
X	011111	-----	10110 01000 0	<b>subfzeo</b>	Subtract From Zero Extended with CA & record OV	Book E
X	011111	-----	10110 01000 1	<b>subfzeo.</b>	Subtract From Zero Extended with CA & record OV & CR	Book E
X	011111	-----	10110 01010 0	<b>addzeo</b>	Add to Zero Extended with CA & record OV	Book E
X	011111	-----	10110 01010 1	<b>addzeo.</b>	Add to Zero Extended with CA & record OV & CR	Book E
X	011111	-----	10110 10101 /	<b>stswi</b>	Store String Word Immediate	Book E
X	011111	-----	10111 01000 0	<b>subfmeo</b>	Subtract From Minus One Extended with CA & record OV	Book E
X	011111	-----	10111 01000 1	<b>subfmeo.</b>	Subtract From Minus One Extended with CA & record OV & CR	Book E
X	011111	-----	10111 01010 0	<b>addmeo</b>	Add to Minus One Extended with CA & record OV	Book E
X	011111	-----	10111 01010 1	<b>addmeo.</b>	Add to Minus One Extended with CA & record OV & CR	Book E
X	011111	-----	10111 01011 0	<b>mullwo</b>	Multiply Low Word & record OV	Book E
X	011111	-----	10111 01011 1	<b>mullwo.</b>	Multiply Low Word & record OV & CR	Book E

Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	10111 10110 /	<b>dcba</b>	Data Cache Block Allocate Indexed	Book E
X	011111	-----	11000 01010 0	<b>addo</b>	Add & record OV	Book E
X	011111	-----	11000 01010 1	<b>addo.</b>	Add & record OV & CR	Book E
X	011111	-----	11000 10010 /	<b>tlbivax</b>	TLB Invalidate Virtual Address Indexed	Book E
X	011111	-----	11000 10110 /	<b>lhbrx</b>	Load Halfword Byte-Reverse Indexed	Book E
X	011111	-----	11000 11000 0	<b>sraw</b>	Shift Right Algebraic Word	Book E
X	011111	-----	11000 11000 1	<b>sraw.</b>	Shift Right Algebraic Word & record CR	Book E
X	011111	-----	11001 11000 0	<b>srawi</b>	Shift Right Algebraic Word Immediate	Book E
X	011111	-----	11001 11000 1	<b>srawi.</b>	Shift Right Algebraic Word Immediate & record CR	Book E
X	011111	-----	11010 10110 /	<b>mbar</b>	Memory Barrier	Book E
X	011111	-----	11100 10010 ?	<b>tlbsx</b>	TLB Search Indexed	Book E
X	011111	-----	11100 10110 /	<b>sthbrx</b>	Store Halfword Byte-Reverse Indexed	Book E
X	011111	-----	11100 11010 0	<b>extsh</b>	Extend Sign Halfword	Book E
X	011111	-----	11100 11010 1	<b>extsh.</b>	Extend Sign Halfword & record CR	Book E
X	011111	-----	11101 10010 /	<b>tlbre</b>	TLB Read Entry	Book E
X	011111	-----	11101 11010 0	<b>extsb</b>	Extend Sign Byte	Book E
X	011111	-----	11101 11010 1	<b>extsb.</b>	Extend Sign Byte & record CR	Book E
X	011111	-----	11110 01011 0	<b>divwuo</b>	Divide Word Unsigned & record OV	Book E
X	011111	-----	11110 01011 1	<b>divwuo.</b>	Divide Word Unsigned & record OV & CR	Book E



Table 266. 32-bit instruction encodings (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	11110 10010 /	<b>tlbwe</b>	TLB Write Entry	Book E
X	011111	-----	11110 10110 /	<b>icbi</b>	Instruction Cache Block Invalidate Indexed	Book E
X	011111	-----	11111 01011 0	<b>divwo</b>	Divide Word & record OV	Book E
X	011111	-----	11111 01011 1	<b>divwo.</b>	Divide Word & record OV & CR	Book E
X	011111	-----	11111 10110 /	<b>dcbz</b>	Data Cache Block set to Zero Indexed	Book E

## 15.4 Instruction index sorted by mnemonic

Table 267 lists all the 16-bit Power\*Embedded instructions, sorted by mnemonic.

Table 267. Instruction index sorted by mnemonic

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )																Mnemonic	Instruction	Page
RR	0	0	0	0	0	1	0	0	y	y	y	y	x	x	x	x	<b>se_add</b>	Add	<a href="#">on page 741</a>
IM5	0	0	1	0	0	0	0	i	i	i	i	i	x	x	x	x	<b>se_addi</b>	Add Immediate	<a href="#">on page 741</a>
RR	0	1	0	0	0	1	1	0	y	y	y	y	x	x	x	x	<b>se_and</b>	AND	<a href="#">on page 743</a>
RR	0	1	0	0	0	1	1	1	y	y	y	y	x	x	x	x	<b>se_and.</b>	AND and Record	<a href="#">on page 743</a>
RR	0	1	0	0	0	1	0	1	y	y	y	y	x	x	x	x	<b>se_andc</b>	AND with Complement	<a href="#">on page 743</a>
IM5	0	0	1	0	1	1	1	i	i	i	i	i	x	x	x	x	<b>se_andi</b>	And Immediate	<a href="#">on page 743</a>
B8	1	1	1	1	0	o	i	i	d	d	d	d	d	d	d	d	<b>se_bc</b>	Branch Conditional	<a href="#">on page 745</a>
IM5	0	1	1	0	0	0	0	i	i	i	i	i	x	x	x	x	<b>se_bclri</b>	Bit Clear Immediate	<a href="#">on page 746</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	<b>se_bctr</b>	Branch to Count Register	<a href="#">on page 746</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	<b>se_bctrl</b>	Branch to Count Register & Link	<a href="#">on page 746</a>
IM5	0	1	1	0	0	0	1	i	i	i	i	i	x	x	x	x	<b>se_bgeni</b>	Bit Generate Immediate	<a href="#">on page 747</a>
B8	1	1	1	1	1	0	0	1	d	d	d	d	d	d	d	d	<b>se_bl</b>	Branch and Link	<a href="#">on page 744</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	<b>se_blr</b>	Branch to Link Register	<a href="#">on page 747</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	<b>se_blrl</b>	Branch to Link Register & Link	<a href="#">on page 747</a>
IM5	0	0	1	0	1	1	0	i	i	i	i	i	x	x	x	x	<b>se_bmaski</b>	Bit Mask Generate Immediate	<a href="#">on page 748</a>
B8	1	1	1	1	1	0	0	0	d	d	d	d	d	d	d	d	<b>se_b</b>	Branch	<a href="#">on page 744</a>
IM5	0	1	1	0	0	1	0	i	i	i	i	i	x	x	x	x	<b>se_bseti</b>	Bit Set Immediate	<a href="#">on page 748</a>
IM5	0	1	1	0	0	1	1	i	i	i	i	i	x	x	x	x	<b>se_btsti</b>	Bit Test Immediate	<a href="#">on page 748</a>
RR	0	0	0	0	1	1	0	0	y	y	y	y	x	x	x	x	<b>se_cmp</b>	Compare	<a href="#">on page 749</a>

Table 267. Instruction index sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )												Mnemonic	Instruction	Page				
IM5	0	0	1	0	1	0	1	i	i	i	i	i	x	x	x	x	<b>se_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
RR	0	0	0	0	1	1	0	1	y	y	y	y	x	x	x	x	<b>se_cmpl</b>	Compare Logical	<a href="#">on page 752</a>
IM5	0	0	1	0	0	0	1	i	i	i	i	i	x	x	x	x	<b>se_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
R	0	0	0	0	0	0	0	0	1	1	0	1	x	x	x	x	<b>se_extsb</b>	Extend Sign Byte	<a href="#">on page 755</a>
R	0	0	0	0	0	0	0	0	1	1	1	1	x	x	x	x	<b>se_extsh</b>	Extend Sign Halfword	<a href="#">on page 755</a>
R	0	0	0	0	0	0	0	0	1	1	0	0	x	x	x	x	<b>se_extzb</b>	Extend with Zeros Byte	<a href="#">on page 756</a>
R	0	0	0	0	0	0	0	0	1	1	1	0	x	x	x	x	<b>se_extzh</b>	Extend with Zeros Halfword	<a href="#">on page 756</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	<b>se_isync</b>	Instruction Synchronize	<a href="#">on page 757</a>
SD4	1	0	0	0	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_lbz</b>	Load Byte and Zero	<a href="#">on page 757</a>
R	0	0	0	0	0	0	0	0	0	1	0	0	x	x	x	x	<b>se_lmww</b>	Load Multiple Word	<a href="#">on page 760</a>
SD4	1	0	1	0	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_lhz</b>	Load Halfword and Zero	<a href="#">on page 759</a>
IM7	0	1	0	0	1	i	i	i	i	i	i	i	x	x	x	x	<b>se_li</b>	Load Immediate	<a href="#">on page 760</a>
SD4	1	1	0	0	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_lwz</b>	Load Word and Zero	<a href="#">on page 761</a>
RR	0	0	0	0	0	0	1	1	y	y	y	y	x	x	x	x	<b>se_mfar</b>	Move from Alternate Register	<a href="#">on page 762</a>
R	0	0	0	0	0	0	0	0	1	0	1	0	x	x	x	x	<b>se_mfctr</b>	Move From Count Register	<a href="#">on page 762</a>
R	0	0	0	0	0	0	0	0	1	0	0	0	x	x	x	x	<b>se_mflr</b>	Move From Link Register	<a href="#">on page 762</a>
RR	0	0	0	0	1	0	1	1	y	y	y	y	x	x	x	x	<b>se_mr</b>	Move Register	<a href="#">on page 762</a>
RR	0	0	0	0	0	0	1	0	y	y	y	y	x	x	x	x	<b>se_mtar</b>	Move to Alternate Register	<a href="#">on page 763</a>
R	0	0	0	0	0	0	0	0	1	0	1	1	x	x	x	x	<b>se_mtctr</b>	Move To Count Register	<a href="#">on page 763</a>
R	0	0	0	0	0	0	0	0	1	0	0	1	x	x	x	x	<b>se_mtlr</b>	Move To Link Register	<a href="#">on page 763</a>

Table 267. Instruction index sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )																Mnemonic	Instruction	Page
RR	0	0	0	0	1	0	1	0	y	y	y	y	x	x	x	x	<b>se_mullw</b>	Multiply Low Word	<a href="#">on page 764</a>
R	0	0	0	0	0	0	0	0	0	0	1	0	x	x	x	x	<b>se_not</b>	NOT	<a href="#">on page 764</a>
RR	0	0	0	0	1	1	1	1	y	y	y	y	x	x	x	x	<b>se_or</b>	OR	<a href="#">on page 765</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	<b>se_rfc</b>	Return From Critical Interrupt	<a href="#">on page 767</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	<b>se_rfdi</b>	Return From Debug Interrupt	<a href="#">on page 704</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	<b>se_rfi</b>	Return From Interrupt	<a href="#">on page 767</a>
RR	0	1	0	0	0	0	0	1	y	y	y	y	x	x	x	x	<b>se_rl</b>	Rotate Left Word	<a href="#">on page 768</a>
IM5	0	1	1	0	1	0	0	i	i	i	i	i	x	x	x	x	<b>se_rlwi</b>	Rotate Left Word Immediate	<a href="#">on page 768</a>
C	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	<b>se_sc</b>	System Call	<a href="#">on page 769</a>
RR	0	1	0	0	0	0	1	1	y	y	y	y	x	x	x	x	<b>se_sl</b>	Shift Left Word	<a href="#">on page 770</a>
IM5	0	1	1	0	1	1	0	i	i	i	i	i	x	x	x	x	<b>se_slwi</b>	Shift Left Word Immediate	<a href="#">on page 770</a>
RR	0	1	0	0	0	0	0	0	y	y	y	y	x	x	x	x	<b>se_s</b>	Shift Right Algebraic Word	<a href="#">on page 771</a>
IM5	0	1	1	0	1	0	1	i	i	i	i	i	x	x	x	x	<b>se_s</b>	Shift Right Algebraic Word Immediate	<a href="#">on page 771</a>
RR	0	1	0	0	0	0	1	0	y	y	y	y	x	x	x	x	<b>se_s</b>	Shift Right Word	<a href="#">on page 772</a>
IM5	0	1	1	0	1	1	1	i	i	i	i	i	x	x	x	x	<b>se_s</b>	Shift Right Word Immediate	<a href="#">on page 772</a>
SD4	1	0	0	1	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_stb</b>	Store Byte	<a href="#">on page 773</a>
SD4	1	0	1	1	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_sth</b>	Store Halfword	<a href="#">on page 773</a>
R	0	0	0	0	0	0	0	0	0	1	0	1	x	x	x	x	<b>se_stmw</b>	Store Multiple Word	<a href="#">on page 774</a>
SD4	1	1	0	1	i	i	i	i	z	z	z	z	x	x	x	x	<b>se_stw</b>	Store Word	<a href="#">on page 775</a>
RR	0	0	0	0	0	1	1	0	y	y	y	y	x	x	x	x	<b>se_sub</b>	Subtract	<a href="#">on page 775</a>

Table 267. Instruction index sorted by mnemonic (continued)

Format	16-Bit Opcodes (Inst <sub>0:15</sub> )												Mnemonic	Instruction	Page				
RR	0	0	0	0	0	1	1	1	y	y	y	y	x	x	x	x	<b>se_sub.</b>	Subtract and Record	<a href="#">on page 775</a>
RR	0	1	0	0	0	1	0	0	y	y	y	y	x	x	x	x	<b>se_subf</b>	Subtract From	<a href="#">on page 776</a>
IM5	0	0	1	0	1	0	0	i	i	i	i	i	x	x	x	x	<b>se_subfic</b>	Subtract From Immediate Carrying	<a href="#">on page 776</a>
IM5	0	0	1	0	0	1	0	i	i	i	i	i	x	x	x	x	<b>se_subi</b>	Subtract Immediate	<a href="#">on page 776</a>
IM5	0	0	1	0	0	1	1	i	i	i	i	i	x	x	x	x	<b>se_subi.</b>	Subtract Immediate and Record	<a href="#">on page 776</a>
RR	0	0	0	0	1	1	1	0	y	y	y	y	x	x	x	x	<b>se_xor</b>	XOR	<a href="#">on page 777</a>

Table 268 sorts 32-bit instructions by mnemonic, ignoring the e\_ prefix.

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	01000 01010 0	<b>add</b>	Add	Book E
X	011111	-----	01000 01010 1	<b>add.</b>	Add & record CR	Book E
X	011111	-----	11000 01010 0	<b>addo</b>	Add & record OV	Book E
X	011111	-----	11000 01010 1	<b>addo.</b>	Add & record OV & CR	Book E
X	011111	-----	00000 01010 0	<b>addc</b>	Add Carrying	Book E
X	011111	-----	00000 01010 1	<b>addc.</b>	Add Carrying & record CR	Book E
X	011111	-----	10000 01010 0	<b>addco</b>	Add Carrying & record OV	Book E
X	011111	-----	10000 01010 1	<b>addco.</b>	Add Carrying & record OV & CR	Book E
X	011111	-----	00100 01010 0	<b>adde</b>	Add Extended with CA	Book E
X	011111	-----	00100 01010 1	<b>adde.</b>	Add Extended with CA & record CR	Book E
X	011111	-----	10100 01010 0	<b>addeo</b>	Add Extended with CA & record OV	Book E
X	011111	-----	10100 01010 1	<b>addeo.</b>	Add Extended with CA & record OV & CR	Book E
D	001111	ttttt aaaaa	iiiiii iiiiii i	<b>e_add16i</b>	Add Immediate	<a href="#">on page 741</a>
SCI8	001110	ttttt aaaaa	10000 FSSii iiiiii i	<b>e_addi</b>	Add Immediate	<a href="#">on page 741</a>
SCI8	001110	ttttt aaaaa	11000 FSSii iiiiii i	<b>e_addi.</b>	Add Immediate and Record	<a href="#">on page 741</a>
SCI8	001110	ttttt aaaaa	10001 FSSii iiiiii i	<b>e_addic</b>	Add Immediate Carrying	<a href="#">on page 743</a>



Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
SCI8	001110	ttttt aaaaa 11001	FSSi i i i i i i	<b>e_addic.</b>	Add Immediate Carrying and Record	<a href="#">on page 743</a>
X	011111	-----	00111 01010 0	<b>addme</b>	Add to Minus One Extended with CA	Book E
X	011111	-----	00111 01010 1	<b>addme.</b>	Add to Minus One Extended with CA & record CR	Book E
X	011111	-----	10111 01010 0	<b>addmeo</b>	Add to Minus One Extended with CA & record OV	Book E
X	011111	-----	10111 01010 1	<b>addmeo.</b>	Add to Minus One Extended with CA & record OV & CR	Book E
X	011111	-----	00110 01010 0	<b>addze</b>	Add to Zero Extended with CA	Book E
X	011111	-----	00110 01010 1	<b>addze.</b>	Add to Zero Extended with CA & record CR	Book E
X	011111	-----	10110 01010 0	<b>addzeo</b>	Add to Zero Extended with CA & record OV	Book E
X	011111	-----	10110 01010 1	<b>addzeo.</b>	Add to Zero Extended with CA & record OV & CR	Book E
X	011111	-----	00000 11100 0	<b>and</b>	AND	Book E
X	011111	-----	00000 11100 1	<b>and.</b>	AND & record CR	Book E
X	011111	-----	00001 11100 0	<b>andc</b>	AND with Complement	Book E
X	011111	-----	00001 11100 1	<b>andc.</b>	AND with Complement & record CR	Book E
SCI8	001110	ttttt aaaaa 10010	FSSi i i i i i i	<b>e_andi</b>	AND Immediate	<a href="#">on page 743</a>
SCI8	001110	ttttt aaaaa 11010	FSSi i i i i i i	<b>e_andi.</b>	AND Immediate and Record	<a href="#">on page 743</a>

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
BD24	011110	0 d d d d d d d d d d d d d d d d	d d d d d d d d d d d d d d d d 0	<b>e_b</b>	Branch	<a href="#">on page 744</a>
BD15	011110	1 0 0 0 o o i i i i d d d d d d	d d d d d d d d d d d d d d d d 0	<b>e_bc</b>	Branch Conditional	<a href="#">on page 745</a>
BD15	011110	1 0 0 0 o o i i i i d d d d d d	d d d d d d d d d d d d d d d d 1	<b>e_bcl</b>	Branch Conditional & Link	<a href="#">on page 745</a>
BD24	011110	0 d d d d d d d d d d d d d d d d	d d d d d d d d d d d d d d d d 1	<b>e_bl</b>	Branch & Link	<a href="#">on page 744</a>
X	011111	- - - - - - - - - - - - - - - -	0 0 0 0 0 0 0 0 0 0 /	<b>cmp</b>	Compare	Book E
SCI8	001110	0 0 0 0 0 a a a a a 0 1 0 b F	F S S i i i i i i i i i i	<b>e_cmpi</b>	Compare Immediate	<a href="#">on page 749</a>
X	011111	- - - - - - - - - - - - - - - -	0 0 0 0 1 0 0 0 0 0 /	<b>cmpl</b>	Compare Logical	Book E
SCI8	001110	0 0 0 0 0 a a a a a 0 1 1 b F	F S S i i i i i i i i i i	<b>e_cmpli</b>	Compare Logical Immediate	<a href="#">on page 752</a>
X	011111	- - - - - - - - - - - - - - - -	0 0 0 0 0 1 1 0 1 0 0	<b>cntlzw</b>	Count Leading Zeros Word	Book E
X	011111	- - - - - - - - - - - - - - - -	0 0 0 0 0 1 1 0 1 0 1	<b>cntlzw.</b>	Count Leading Zeros Word & record CR	Book E
X	011111	- - - - - - - - - - - - - - - -	1 0 1 1 1 1 0 1 1 0 /	<b>dcba</b>	Data Cache Block Allocate Indexed	Book E
X	011111	- - - - - - - - - - - - - - - -	0 0 0 1 0 1 0 1 1 0 /	<b>dcbf</b>	Data Cache Block Flush Indexed	Book E
X	011111	- - - - - - - - - - - - - - - -	0 1 1 1 0 1 0 1 1 0 /	<b>dcbi</b>	Data Cache Block Invalidate Indexed	Book E
X	011111	- - - - - - - - - - - - - - - -	0 0 0 0 1 1 0 1 1 0 /	<b>dcbst</b>	Data Cache Block Store Indexed	Book E
X	011111	- - - - - - - - - - - - - - - -	0 0 1 1 1 1 0 1 1 0 /	<b>dcbtst</b>	Data Cache Block Touch for Store Indexed	Book E
X	011111	- - - - - - - - - - - - - - - -	0 1 0 0 0 1 0 1 1 0 /	<b>dcbt</b>	Data Cache Block Touch Indexed	Book E





Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1	-----	1 1 1 1 1 1 0 1 1 1 0 /	<b>dcbz</b>	Data Cache Block set to Zero Indexed	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 1 0 1 0 1 1 1 0	<b>divw</b>	Divide Word	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 1 0 1 0 1 1 1 1	<b>divw.</b>	Divide Word & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 1 1 1 0 1 0 1 1 1 0	<b>divwo</b>	Divide Word & record OV	Book E
X	0 1 1 1 1 1	-----	1 1 1 1 1 0 1 0 1 1 1 1	<b>divwo.</b>	Divide Word & record OV & CR	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 0 0 1 0 1 1 1 0	<b>divwu</b>	Divide Word Unsigned	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 0 0 1 0 1 1 1 1	<b>divwu.</b>	Divide Word Unsigned & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 1 1 0 0 1 0 1 1 1 0	<b>divwuo</b>	Divide Word Unsigned & record OV	Book E
X	0 1 1 1 1 1	-----	1 1 1 1 0 0 1 0 1 1 1 1	<b>divwuo.</b>	Divide Word Unsigned & record OV & CR	Book E
X	0 1 1 1 1 1	-----	0 1 0 0 0 1 1 1 0 0 0 0	<b>eqv</b>	Equivalent	Book E
X	0 1 1 1 1 1	-----	0 1 0 0 0 1 1 1 0 0 0 1	<b>eqv.</b>	Equivalent & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 1 0 1 1 1 0 1 0 0 0	<b>extsb</b>	Extend Sign Byte	Book E
X	0 1 1 1 1 1	-----	1 1 1 0 1 1 1 0 1 0 0 1	<b>extsb.</b>	Extend Sign Byte & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 1 0 0 1 1 0 1 0 0 0	<b>extsh</b>	Extend Sign Halfword	Book E
X	0 1 1 1 1 1	-----	1 1 1 0 0 1 1 0 1 0 0 1	<b>extsh.</b>	Extend Sign Halfword & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 1 1 0 1 0 1 1 1 0 /	<b>icbi</b>	Instruction Cache Block Invalidate Indexed	Book E
X	0 1 1 1 1 1	-----	0 0 0 0 0 1 0 1 1 1 0 /	<b>icbt</b>	Instruction Cache Block Touch Indexed	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0111111	-----	----- 011111 /	<b>isel</b>	Integer Select	Book E
D14	0011100	ttttt aaaaa 00ddd	dddd dddd d	<b>e_lbz</b>	Load Byte & Zero	<a href="#">on page 757</a>
X	0111111	-----	00010 10111 /	<b>lbzx</b>	Load Byte & Zero Indexed	Book E
D8	0011110	ttttt aaaaa 00000	000dd dddd d	<b>e_lbzu</b>	Load Byte & Zero with Update	<a href="#">on page 757</a>
X	0111111	-----	00011 10111 /	<b>lbzux</b>	Load Byte & Zero with Update Indexed	Book E
D14	0011100	ttttt aaaaa 11ddd	dddd dddd d	<b>e_ld</b>	Load Doubleword & Zero (reserved for 64b GPR)	
D8	0011110	ttttt aaaaa 00000	011dd dddd d	<b>e_ldu</b>	Load Doubleword with Update (reserved for 64b GPR)	
D14	0011100	ttttt aaaaa 01ddd	dddd dddd d	<b>e_lhz</b>	Load Halfword & Zero	<a href="#">on page 759</a>
D8	0011110	ttttt aaaaa 00000	001dd dddd d	<b>e_lhzu</b>	Load Halfword & Zero with Update	<a href="#">on page 759</a>
X	0111111	-----	01000 10111 /	<b>lhzx</b>	Load Halfword & Zero Indexed	Book E
X	0111111	-----	01001 10111 /	<b>lhzux</b>	Load Halfword & Zero with Update Indexed	Book E
D8	0011110	ttttt aaaaa 00001	010dd dddd d	<b>e_lha</b>	Load Halfword Algebraic	<a href="#">on page 758</a>
X	0111111	-----	01010 10111 /	<b>lhax</b>	Load Halfword Algebraic Indexed	Book E
D8	0011110	ttttt aaaaa 00001	011dd dddd d	<b>e_lhau</b>	Load Halfword Algebraic with Update	<a href="#">on page 758</a>
X	0111111	-----	01011 10111 /	<b>lhaux</b>	Load Halfword Algebraic with Update Indexed	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	11000 10110 /	<b>lhbrx</b>	Load Halfword Byte-Reverse Indexed	Book E
LI20	011100	ttttt 0iiii iiii i	iiii iiii i	<b>e_li</b>	Load Immediate	<a href="#">on page 760</a>
LI20	011100	ttttt 1iiii iiii i	iiii iiii i	<b>e_lis</b>	Load Immediate Shifted	<a href="#">on page 760</a>
D8	001110	ttttt aaaaa 00001	000dd dddd d	<b>e_lmw</b>	Load Multiple Word	<a href="#">on page 760</a>
X	011111	-----	00000 10100 /	<b>lwarx</b>	Load Word & Reserve Indexed	Book E
X	011111	-----	10000 10110 /	<b>lwbrx</b>	Load Word Byte-Reverse Indexed	Book E
D14	001100	ttttt aaaaa 10ddd	dddd dddd d	<b>e_lwz</b>	Load Word & Zero	<a href="#">on page 761</a>
D8	001110	ttttt aaaaa 00000	010dd dddd d	<b>e_lwzu</b>	Load Word & Zero with Update	<a href="#">on page 761</a>
X	011111	-----	00000 10111 /	<b>lwzx</b>	Load Word & Zero Indexed	Book E
X	011111	-----	00001 10111 /	<b>lwzux</b>	Load Word & Zero with Update Indexed	Book E
X	011111	-----	11010 10110 /	<b>mbar</b>	Memory Barrier	Book E
X	011111	-----	01000 10011 /	<b>mfapidi</b>	Move From APID Indirect	Book E
X	011111	-----	00000 10011 /	<b>mfcr</b>	Move From Condition Register	Book E
XFX	011111	-----	01010 00011 /	<b>mfocr</b>	Move From Device Control Register	Book E
X	011111	-----	00010 10011 /	<b>mfmsr</b>	Move From Machine State Register	Book E
XFX	011111	-----	01010 10011 /	<b>mfspir</b>	Move From Special Purpose Register	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1	-----	1 0 0 1 0 1 0 1 1 0 /	<b>msync</b>	Memory Synchronize	Book E
AFX	0 1 1 1 1 1	-----	0 0 1 0 0 1 0 0 0 0 /	<b>mtcrf</b>	Move To Condition Register Fields	Book E
X	0 1 1 1 1 1	-----	1 0 0 0 0 0 0 0 0 0 /	<b>mcrxr</b>	Move to Condition Register from XER	Book E
AFX	0 1 1 1 1 1	-----	0 1 1 1 0 0 0 0 1 1 /	<b>mtdcr</b>	Move To Device Control Register	Book E
X	0 1 1 1 1 1	-----	0 0 1 0 0 1 0 0 1 0 /	<b>mtmsr</b>	Move To Machine State Register	Book E
AFX	0 1 1 1 1 1	-----	0 1 1 1 0 1 0 0 1 1 /	<b>mtspr</b>	Move To Special Purpose Register	Book E
X	0 1 1 1 1 1	-----	/ 0 0 1 0 0 1 0 1 1 0	<b>mulhw</b>	Multiply High Word	Book E
X	0 1 1 1 1 1	-----	/ 0 0 1 0 0 1 0 1 1 1	<b>mulhw.</b>	Multiply High Word & record CR	Book E
X	0 1 1 1 1 1	-----	/ 0 0 0 0 0 1 0 1 1 0	<b>mulhwu</b>	Multiply High Word Unsigned	Book E
X	0 1 1 1 1 1	-----	/ 0 0 0 0 0 1 0 1 1 1	<b>mulhwu.</b>	Multiply High Word Unsigned & record CR	Book E
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 0 1 1 0 F S S i i i i i i i	<b>e_mulli</b>	Multiply Low Immediate	<a href="#">on page 763</a>
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 1 1 1 0 F S S i i i i i i i	<b>e_mulli.</b>	Multiply Low Immediate and Record	<a href="#">on page 763</a>
X	0 1 1 1 1 1	-----	0 0 1 1 1 0 1 0 1 1 0	<b>mullw</b>	Multiply Low Word	Book E
X	0 1 1 1 1 1	-----	0 0 1 1 1 0 1 0 1 1 1	<b>mullw.</b>	Multiply Low Word & record CR	Book E
X	0 1 1 1 1 1	-----	1 0 1 1 1 0 1 0 1 1 0	<b>mullwo</b>	Multiply Low Word & record OV	Book E
X	0 1 1 1 1 1	-----	1 0 1 1 1 0 1 0 1 1 1	<b>mullwo.</b>	Multiply Low Word & record OV & CR	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 0 1 1 1 0 0 0	<b>nand</b>	NAND	Book E
X	0 1 1 1 1 1	-----	0 1 1 1 0 1 1 1 0 0 1	<b>nand.</b>	NAND & record CR	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1	-----	0 0 0 1 1 0 1 0 0 0 0	<b>neg</b>	Negate	Book E
X	0 1 1 1 1 1	-----	0 0 0 1 1 0 1 0 0 0 1	<b>neg.</b>	Negate & record CR	Book E
X	0 1 1 1 1 1	-----	1 0 0 1 1 0 1 0 0 0 0	<b>nego</b>	Negate & record OV	Book E
X	0 1 1 1 1 1	-----	1 0 0 1 1 0 1 0 0 0 1	<b>nego.</b>	Negate & record OV & record CR	Book E
X	0 1 1 1 1 1	-----	0 0 0 1 1 1 1 1 0 0 0	<b>nor</b>	NOR	Book E
X	0 1 1 1 1 1	-----	0 0 0 1 1 1 1 1 0 0 1	<b>nor.</b>	NOR & record CR	Book E
X	0 1 1 1 1 1	-----	0 1 1 0 1 1 1 1 0 0 0	<b>or</b>	OR	Book E
X	0 1 1 1 1 1	-----	0 1 1 0 1 1 1 1 0 0 1	<b>or.</b>	OR & record CR	Book E
X	0 1 1 1 1 1	-----	0 1 1 0 0 1 1 1 0 0 0	<b>orc</b>	OR with Complement	Book E
X	0 1 1 1 1 1	-----	0 1 1 0 0 1 1 1 0 0 1	<b>orc.</b>	OR with Complement & record CR	Book E
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 0 0 1 1 F S S i i i i i i i	<b>e_ori</b>	OR Immediate	<a href="#">on page 765</a>
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 1 0 1 1 F S S i i i i i i i	<b>e_ori.</b>	OR Immediate and Record	<a href="#">on page 768</a>
	0 1 1 1 0 1	S S S S S a a a a a h h h h h	b b b b b e e e e e 1	<b>e_rlwinm.</b>	Rotate Left Word Immed then AND with Mask & record CR	<a href="#">on page 769</a>
	0 1 1 1 0 1	S S S S S a a a a a h h h h h	b b b b b e e e e e 0	<b>e_rlwimi.</b>	Rotate Left Word Immed then Mask Insert & record CR	<a href="#">on page 768</a>
X	0 1 1 1 1 1	-----	0 0 0 0 0 1 1 0 0 0 0	<b>slw</b>	Shift Left Word	Book E
X	0 1 1 1 1 1	-----	0 0 0 0 0 1 1 0 0 0 1	<b>slw.</b>	Shift Left Word & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 0 0 0 1 1 0 0 0 0	<b>sraw</b>	Shift Right Algebraic Word	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1	-----	1 1 0 0 0 1 1 0 0 0 1	<b>sraw.</b>	Shift Right Algebraic Word & record CR	Book E
X	0 1 1 1 1 1	-----	1 1 0 0 1 1 1 0 0 0 0	<b>srawi</b>	Shift Right Algebraic Word Immediate	Book E
X	0 1 1 1 1 1	-----	1 1 0 0 1 1 1 0 0 0 1	<b>srawi.</b>	Shift Right Algebraic Word Immediate & record CR	Book E
X	0 1 1 1 1 1	-----	1 0 0 0 0 1 1 0 0 0 0	<b>srw</b>	Shift Right Word	Book E
X	0 1 1 1 1 1	-----	1 0 0 0 0 1 1 0 0 0 1	<b>srw.</b>	Shift Right Word & record CR	Book E
D14	0 0 1 1 0 1	t t t t t a a a a a 0 0 d d d	d d d d d d d d d d d	<b>e_stb</b>	Store Byte	<a href="#">on page 773</a>
X	0 1 1 1 1 1	-----	0 0 1 1 0 1 0 1 1 1 /	<b>stbx</b>	Store Byte Indexed	Book E
D8	0 0 1 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 0 0 d d d d d d d d	<b>e_stbu</b>	Store Byte with Update	<a href="#">on page 773</a>
X	0 1 1 1 1 1	-----	0 0 1 1 1 1 0 1 1 1 /	<b>stbux</b>	Store Byte with Update Indexed	<a href="#">on page 773</a>
D14	0 0 1 1 0 1	t t t t t a a a a a 1 1 d d d	d d d d d d d d d d d	<b>e_std</b>	Store Doubleword (reserved for 64b GPR)	Book E
D8	0 0 1 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 1 1 d d d d d d d d	<b>e_stdu</b>	Store Doubleword with Update (reserved for 64b GPR)	
D14	0 0 1 1 0 1	t t t t t a a a a a 0 1 d d d	d d d d d d d d d d d	<b>e_sth</b>	Store Halfword	<a href="#">on page 773</a>
X	0 1 1 1 1 1	-----	1 1 1 0 0 1 0 1 1 0 /	<b>sthbrx</b>	Store Halfword Byte-Reverse Indexed	Book E
X	0 1 1 1 1 1	-----	0 1 1 0 0 1 0 1 1 1 /	<b>sthx</b>	Store Halfword Indexed	Book E
D8	0 0 1 1 1 0	t t t t t a a a a a 0 0 0 0 0	1 0 1 d d d d d d d d	<b>e_sthu</b>	Store Halfword with Update	<a href="#">on page 773</a>



Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	01101 10111 /	<b>sthux</b>	Store Halfword with Update Indexed	Book E
D8	001110	ttttt aaaaa 00001	001dd ddddd d	<b>e_stmw</b>	Store Multiple Word	Book E
D14	001101	ttttt aaaaa 10ddd	dddd ddddd d	<b>e_stw</b>	Store Word	<a href="#">on page 775</a>
X	011111	-----	10100 10110 /	<b>stwbrx</b>	Store Word Byte-Reverse Indexed	Book E
X	011111	-----	00100 10110 1	<b>stwcx.</b>	Store Word Conditional Indexed & record CR	Book E
D8	001110	ttttt aaaaa 00000	110dd ddddd d	<b>e_stwu</b>	Store Word with Update	<a href="#">on page 775</a>
X	011111	-----	00101 10111 /	<b>stwux</b>	Store Word with Update Indexed	Book E
X	011111	-----	00100 10111 /	<b>stwx</b>	Store Word Indexed	Book E
X	011111	-----	00001 01000 0	<b>subf</b>	Subtract From	Book E
X	011111	-----	00001 01000 1	<b>subf.</b>	Subtract From & record CR	Book E
X	011111	-----	00000 01000 0	<b>subfc</b>	Subtract From Carrying	Book E
X	011111	-----	00000 01000 1	<b>subfc.</b>	Subtract From Carrying & record CR	Book E
X	011111	-----	10000 01000 0	<b>subfco</b>	Subtract From Carrying & record OV	Book E
X	011111	-----	10000 01000 1	<b>subfco.</b>	Subtract From Carrying & record OV & CR	Book E
X	011111	-----	10001 01000 0	<b>subfo</b>	Subtract From & record OV	Book E
X	011111	-----	10001 01000 1	<b>subfo.</b>	Subtract From & record OV & CR	Book E
X	011111	-----	00100 01000 0	<b>subfe</b>	Subtract From Extended with CA	Book E

Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	011111	-----	00100 01000 1	<b>subfe.</b>	Subtract From Extended with CA & record CR	Book E
X	011111	-----	10100 01000 0	<b>subfeo</b>	Subtract From Extended with CA & record OV	Book E
X	011111	-----	10100 01000 1	<b>subfeo.</b>	Subtract From Extended with CA & record OV & CR	Book E
SCI8	001110	ttttt aaaaa 10100	FSSi i i i i i i	<b>e_subfic</b>	Subtract from Immediate Carrying	<a href="#">on page 776</a>
SCI8	001110	ttttt aaaaa 11100	FSSi i i i i i i	<b>e_subfic.</b>	Subtract from Immediate Carrying and Record	<a href="#">on page 776</a>
X	011111	-----	00111 01000 0	<b>subfme</b>	Subtract From Minus One Extended with CA	Book E
X	011111	-----	00111 01000 1	<b>subfme.</b>	Subtract From Minus One Extended with CA & record CR	Book E
X	011111	-----	10111 01000 0	<b>subfmeo</b>	Subtract From Minus One Extended with CA & record OV	Book E
X	011111	-----	10111 01000 1	<b>subfmeo.</b>	Subtract From Minus One Extended with CA & record OV & CR	Book E
X	011111	-----	00110 01000 0	<b>subfze</b>	Subtract From Zero Extended with CA	Book E
X	011111	-----	00110 01000 1	<b>subfze.</b>	Subtract From Zero Extended with CA & record CR	Book E
X	011111	-----	10110 01000 0	<b>subfzeo</b>	Subtract From Zero Extended with CA & record OV	Book E
X	011111	-----	10110 01000 1	<b>subfzeo.</b>	Subtract From Zero Extended with CA & record OV & CR	Book E



Table 268. 32-bit instructions by mnemonic (ignoring the e\_ prefix) (continued)

Format	Opcode			Mnemonic	Instruction	Page
	Primary (Inst <sub>0:5</sub> )	Intermediate (Inst <sub>6:20</sub> )	Extended (Inst <sub>21:31</sub> )			
X	0 1 1 1 1 1	- - - - -	1 1 0 0 0 1 0 0 1 0 /	<b>tlbivax</b>	TLB Invalidate Virtual Address Indexed	Book E
X	0 1 1 1 1 1	- - - - -	1 1 1 0 1 1 0 0 1 0 /	<b>tlbre</b>	TLB Read Entry	Book E
X	0 1 1 1 1 1	- - - - -	1 1 1 0 0 1 0 0 1 0 ?	<b>tlbsx</b>	TLB Search Indexed	Book E
X	0 1 1 1 1 1	- - - - -	1 0 0 0 1 1 0 1 1 0 /	<b>tlbsync</b>	TLB Synchronize	Book E
X	0 1 1 1 1 1	- - - - -	1 1 1 1 0 1 0 0 1 0 /	<b>tlbwe</b>	TLB Write Entry	Book E
X	0 1 1 1 1 1	- - - - -	0 0 0 0 0 0 0 1 0 0 /	<b>tw</b>	Trap Word	Book E
X	0 1 1 1 1 1	- - - - -	0 0 1 0 0 0 0 0 1 1 /	<b>wrtee</b>	Write External Enable	Book E
X	0 1 1 1 1 1	- - - - -	0 0 1 0 1 0 0 0 1 1 /	<b>wrteei</b>	Write External Enable Immediate	Book E
X	0 1 1 1 1 1	- - - - -	0 1 0 0 1 1 1 1 0 0 0	<b>xor</b>	XOR	Book E
X	0 1 1 1 1 1	- - - - -	0 1 0 0 1 1 1 1 0 0 1	<b>xor.</b>	XOR & record CR	Book E
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 0 1 1 1 F S S i i i i i i i i	<b>e_xori</b>	XOR Immediate	<a href="#">on page 777</a>
SCI8	0 0 1 1 1 0	t t t t t a a a a a	1 1 1 1 1 F S S i i i i i i i i	<b>e_xori.</b>	XOR Immediate and Record	<a href="#">on page 777</a>

## Appendix A Instruction set listings

This appendix lists the instructions by both mnemonic and opcode, and includes a quick reference table with general information, such as the architecture level, privilege level, form, and whether the instruction is optional. The tables in the chapter are organized as follows:

- [A.1: Instructions sorted by mnemonic \(decimal and hexadecimal\)](#)
- [A.2: Instructions sorted by primary opcodes \(decimal and hexadecimal\)](#)
- [A.3: Instructions sorted by mnemonic \(binary\)](#)
- [A.4: Instructions sorted by opcode \(binary\)](#)
- [A.5: Instruction set legend](#)

Note that this appendix does not include instructions defined by the VLE extension. These instructions are listed in [Chapter 15: VLE instruction index](#).

### A.1 Instructions sorted by mnemonic (decimal and hexadecimal)

[Table 269](#) lists instructions in alphabetical order by mnemonic, showing decimal and hexadecimal values of the primary opcode (0–5) and binary values of the secondary opcode (21–31). This list also includes simplified mnemonics and their equivalents using standard mnemonics.



**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
<b>add</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	0	1	0	0	X	<b>add</b>		
<b>add.</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	0	1	0	1	X	<b>add.</b>		
<b>addc</b>	31 (0x1F)					rD					rA					rB					0	0	0	0	0	0	1	0	1	0	0	X	<b>addc</b>		
<b>addc.</b>	31 (0x1F)					rD					rA					rB					0	0	0	0	0	0	1	0	1	0	1	X	<b>addc.</b>		
<b>addco</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	1	0	1	0	0	X	<b>addco</b>		
<b>addco.</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	1	0	1	0	1	X	<b>addco.</b>		
<b>adde</b>	31 (0x1F)					rD					rA					rB					0	0	1	0	0	0	1	0	1	0	0	X	<b>adde</b>		
<b>adde.</b>	31 (0x1F)					rD					rA					rB					0	0	1	0	0	0	1	0	1	0	1	X	<b>adde.</b>		
<b>addeo</b>	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	1	0	1	0	0	X	<b>addeo</b>		
<b>addeo.</b>	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	1	0	1	0	1	X	<b>addeo.</b>		
<b>addi</b>	14 (0x0E)					rD					rA					SIMM										D	<b>addi</b>								
<b>addic</b>	12 (0x0C)					rD					rA					SIMM										D	<b>addic</b>								
<b>addic.</b>	13 (0x0D)					rD					rA					SIMM										D	<b>addic.</b>								
<b>addis</b>	15 (0x0F)					rD					rA					SIMM										D	<b>addis</b>								
<b>addme</b>	31 (0x1F)					rD					rA					///					0	0	1	1	1	0	1	0	1	0	1	0	0	X	<b>addme</b>
<b>addme.</b>	31 (0x1F)					rD					rA					///					0	0	1	1	1	0	1	0	1	0	1	X	<b>addme.</b>		
<b>addmeo</b>	31 (0x1F)					rD					rA					///					1	0	1	1	1	0	1	0	1	0	0	X	<b>addmeo</b>		
<b>addmeo.</b>	31 (0x1F)					rD					rA					///					1	0	1	1	1	0	1	0	1	0	1	X	<b>addmeo.</b>		
<b>addo</b>	31 (0x1F)					rD					rA					rB					1	1	0	0	0	0	1	0	1	0	0	X	<b>addo</b>		
<b>addo.</b>	31 (0x1F)					rD					rA					rB					1	1	0	0	0	0	1	0	1	0	1	X	<b>addo.</b>		
<b>addze</b>	31 (0x1F)					rD					rA					///					0	0	1	1	0	0	1	0	1	0	0	X	<b>addze</b>		
<b>addze.</b>	31 (0x1F)					rD					rA					///					0	0	1	1	0	0	1	0	1	0	1	X	<b>addze.</b>		
<b>addzeo</b>	31 (0x1F)					rD					rA					///					1	0	1	1	0	0	1	0	1	0	0	X	<b>addzeo</b>		
<b>addzeo.</b>	31 (0x1F)					rD					rA					///					1	0	1	1	0	0	1	0	1	0	1	X	<b>addzeo.</b>		



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>and</b>	31 (0x1F)					rS					rA					rB					0	0	0	0	0	0	1	1	1	0	0	0	X	<b>and</b>
<b>and.</b>	31 (0x1F)					rS					rA					rB					0	0	0	0	0	0	1	1	1	0	0	1	X	<b>and.</b>
<b>andc</b>	31 (0x1F)					rS					rA					rB					0	0	0	0	1	1	1	1	0	0	0	X	<b>andc</b>	
<b>andc.</b>	31 (0x1F)					rS					rA					rB					0	0	0	0	1	1	1	1	0	0	1	X	<b>andc.</b>	
<b>andi.</b>	28 (0x1C)					rS					rA					UIMM										D	<b>andi.</b>							
<b>andis.</b>	29 (0x1D)					rS					rA					UIMM										D	<b>andis.</b>							
<b>b</b>	18 (0x12)					LI										0	0	I	<b>b</b>															
<b>ba</b>	18 (0x12)					LI										1	0	I	<b>ba</b>															
<b>bc</b>	16 (0x10)					BO					BI					BD					0	0	B	<b>bc</b>										
<b>bca</b>	16 (0x10)					BO					BI					BD					1	0	B	<b>bca</b>										
<b>bcctr</b>	19 (0x13)					BO					BI					///	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	XL	<b>bcctr</b>	
<b>bcctrl</b>	19 (0x13)					BO					BI					///	1	0	0	0	0	0	1	0	0	0	0	0	1	XL	<b>bcctrl</b>			
<b>bcl</b>	16 (0x10)					BO					BI					BD					0	1	B	<b>bcl</b>										
<b>bcla</b>	16 (0x10)					BO					BI					BD					1	1	B	<b>bcla</b>										
<b>bclr</b>	19 (0x13)					BO					BI					///	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	XL	<b>bclr</b>	
<b>bclrl</b>	19 (0x13)					BO					BI					///	0	0	0	0	0	0	1	0	0	0	0	0	1	XL	<b>bclrl</b>			
<b>bctr</b>	<b>bctr</b> <sup>(1)</sup>					equivalent to <b>bcctr 20,0</b>																<b>bctr</b>												
<b>bctrl</b>	<b>bctrl</b> <sup>(1)</sup>					equivalent to <b>bcctrl 20,0</b>																<b>bctrl</b>												
<b>bdnz</b>	<b>bdnz target</b> <sup>(1)</sup>					equivalent to										<b>bc 16,0,target</b>										<b>bdnz</b>								
<b>bdnza</b>	<b>bdnza target</b> <sup>(1)</sup>					equivalent to										<b>bca 16,0,target</b>										<b>bdnza</b>								
<b>bdnzf</b>	<b>bdnzf BI,target</b>					equivalent to										<b>bc 0,BI,target</b>										<b>bdnzf</b>								
<b>bdnzfa</b>	<b>bdnzfa BI,target</b>					equivalent to										<b>bca 0,BI,target</b>										<b>bdnzfa</b>								
<b>bdnzfl</b>	<b>bdnzfl BI,target</b>					equivalent to										<b>bcl 0,BI,target</b>										<b>bdnzfl</b>								
<b>bdnzfla</b>	<b>bdnzfla BI,target</b>					equivalent to										<b>bcla 0,BI,target</b>										<b>bdnzfla</b>								
<b>bdnzflr</b>	<b>bdnzflr BI</b>					equivalent to										<b>bclr 0,BI</b>										<b>bdnzflr</b>								



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>bdnzflrl</b>	<b>bdnzflrl BI</b>																equivalent to										<b>bclrl 0,BI</b>		<b>bdnzflrl</b>					
<b>bdnzl</b>	<b>bdnzl target</b> <sup>(1)</sup>																equivalent to										<b>bcl 16,0,target</b>		<b>bdnzl</b>					
<b>bdnzla</b>	<b>bdnzla target</b> <sup>(1)</sup>																equivalent to										<b>bcla 16,0,target</b>		<b>bdnzla</b>					
<b>bdnzlr</b>	<b>bdnzlr BI</b>																equivalent to										<b>bclr 16,BI</b>		<b>bdnzlr</b>					
<b>bdnzlrl</b>	<b>bdnzlrl</b> <sup>(1)</sup>																equivalent to										<b>bclrl 16,0</b>		<b>bdnzlrl</b>					
<b>bdnzt</b>	<b>bdnzt BI,target</b>																equivalent to										<b>bc 8,BI,target</b>		<b>bdnzt</b>					
<b>bdnzta</b>	<b>bdnzta BI,target</b>																equivalent to										<b>bca 8,BI,target</b>		<b>bdnzta</b>					
<b>bdnztl</b>	<b>bdnztl BI,target</b>																equivalent to										<b>bcl 8,0,target</b>		<b>bdnztl</b>					
<b>bdnztla</b>	<b>bdnztla BI,target</b>																equivalent to										<b>bcla 8,BI,target</b>		<b>bdnztla</b>					
<b>bdnztlr</b>	<b>bdnztlr BI</b>																equivalent to										<b>bclr 8,BI</b>		<b>bdnztlr</b>					
<b>bdnztlrl</b>	<b>bdnztlrl BI</b>																equivalent to										<b>bclrl 8,BI</b>		<b>bdnztlrl</b>					
<b>bdz</b>	<b>bdz target</b> <sup>(1)</sup>																equivalent to										<b>bc 18,0,target</b>		<b>bdz</b>					
<b>bdza</b>	<b>bdza target</b> <sup>(1)</sup>																equivalent to										<b>bca 18,0,target</b>		<b>bdza</b>					
<b>bdzf</b>	<b>bdzf BI,target</b>																equivalent to										<b>bc 2,BI,target</b>		<b>bdzf</b>					
<b>bdzfa</b>	<b>bdzfa BI,target</b>																equivalent to										<b>bca 2,BI,target</b>		<b>bdzfa</b>					
<b>bdzfl</b>	<b>bdzfl BI,target</b>																equivalent to										<b>bcl 2,BI,target</b>		<b>bdzfl</b>					
<b>bdzfla</b>	<b>bdzfla BI,target</b>																equivalent to										<b>bcla 2,BI,target</b>		<b>bdzfla</b>					
<b>bdzflr</b>	<b>bdzflr BI</b>																equivalent to										<b>bclr 2,BI</b>		<b>bdzflr</b>					
<b>bdzflrl</b>	<b>bdzflrl BI</b>																equivalent to										<b>bclrl 2,BI</b>		<b>bdzflrl</b>					
<b>bdzl</b>	<b>bdzl target</b> <sup>(1)</sup>																equivalent to										<b>bcl 18,BI,target</b>		<b>bdzl</b>					
<b>bdzla</b>	<b>bdzla target</b> <sup>(1)</sup>																equivalent to										<b>bcla 18,BI,target</b>		<b>bdzla</b>					
<b>bdzlr</b>	<b>bdzlr</b> <sup>(1)</sup>																equivalent to										<b>bclr 18,0</b>		<b>bdzlr</b>					
<b>bdzlrl</b>	<b>bdzlrl</b> <sup>(1)</sup>																equivalent to										<b>bclrl 18,0</b>		<b>bdzlrl</b>					
<b>bdzt</b>	<b>bdzt BI,target</b>																equivalent to										<b>bc 10,BI,target</b>		<b>bdzt</b>					



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>bdzta</b>	bdzta BI,target		equivalent to										bca 10,BI,target										<b>bdzta</b>											
<b>bdztl</b>	bdztl BI,target		equivalent to										bcl 10,BI,target										<b>bdztl</b>											
<b>bdztla</b>	bdztla BI,target		equivalent to										bcla 10,BI,target										<b>bdztla</b>											
<b>bdztlrl</b>	bdztlrl BI		equivalent to										bclrl 10, BI										<b>bdztlrl</b>											
<b>beq</b>	beq crS,target		equivalent to										bc 12,BI <sup>(2)</sup> ,target										<b>beq</b>											
<b>beqa</b>	beqa crS,target		equivalent to										bca 12,BI <sup>(2)</sup> ,target										<b>beqa</b>											
<b>beqctr</b>	beqctr crS,target		equivalent to										bcctr 12,BI <sup>(2)</sup> ,target										<b>beqctr</b>											
<b>beqctrl</b>	beqctrl crS,target		equivalent to										bcctrl 12,BI <sup>(2)</sup> ,target										<b>beqctrl</b>											
<b>beql</b>	beql crS,target		equivalent to										bcl 12,BI <sup>(2)</sup> ,target										<b>beql</b>											
<b>beqla</b>	beqla crS,target		equivalent to										bcla 12,BI <sup>(2)</sup> ,target										<b>beqla</b>											
<b>beqlr</b>	beqlr crS,target		equivalent to										bclr 12,BI <sup>(2)</sup> ,target										<b>beqlr</b>											
<b>beqlrl</b>	beqlrl crS,target		equivalent to										bclrl 12,BI <sup>(2)</sup> ,target										<b>beqlrl</b>											
<b>bf</b>	bf BI,target		equivalent to										bc 4,BI,target										<b>bf</b>											
<b>bfa</b>	bfa BI,target		equivalent to										bca 4,BI,target										<b>bfa</b>											
<b>bfctr</b>	bfctr BI		equivalent to										bcctr 4,BI										<b>bfctr</b>											
<b>bfctrl</b>	bfctrl BI		equivalent to										bcctrl 4, BI										<b>bfctrl</b>											
<b>bfl</b>	bfl BI,target		equivalent to										bcl 4,BI,target										<b>bfl</b>											
<b>bfla</b>	bfla BI,target		equivalent to										bcla 4,BI,target										<b>bfla</b>											
<b>bflr</b>	bflr BI		equivalent to										bclr 4,BI										<b>bflr</b>											
<b>bflrl</b>	bflrl BI		equivalent to										bclrl 4,BI										<b>bflrl</b>											
<b>bge</b>	bge crS,target		equivalent to										bc 4,BI <sup>(3)</sup> ,target										<b>bge</b>											
<b>bgea</b>	bgea crS,target		equivalent to										bca 4,BI <sup>(3)</sup> ,target										<b>bgea</b>											
<b>bgectr</b>	bgectr crS,target		equivalent to										bcctr 4,BI <sup>(3)</sup> ,target										<b>bgectr</b>											
<b>bgectrl</b>	bgectrl crS,target		equivalent to										bcctrl 4,BI <sup>(3)</sup> ,target										<b>bgectrl</b>											
<b>bgel</b>	bgel crS,target		equivalent to										bcl 4,BI <sup>(3)</sup> ,target										<b>bgel</b>											



**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>bgela</b>	bgela crS,target		equivalent to		bcla 4,BI <sup>(3)</sup> ,target																						<b>bgela</b>							
<b>bgelr</b>	bgelr crS,target		equivalent to		bclr 4,BI <sup>(3)</sup> ,target																						<b>bgelr</b>							
<b>bgelrl</b>	bgelrl crS,target		equivalent to		bclrl 4,BI <sup>(3)</sup> ,target																						<b>bgelrl</b>							
<b>bgt</b>	bgt crS,target		equivalent to		bc 12,BI <sup>(4)</sup> ,target																						<b>bgt</b>							
<b>bgta</b>	bgta crS,target		equivalent to		bca 12,BI <sup>(4)</sup> ,target																						<b>bgta</b>							
<b>bgtctr</b>	bgtctr crS,target		equivalent to		bcctr 12,BI <sup>(4)</sup> ,target																						<b>bgtctr</b>							
<b>bgtctrl</b>	bgtctrl crS,target		equivalent to		bcctrl 12,BI <sup>(4)</sup> ,target																						<b>bgtctrl</b>							
<b>bgtl</b>	bgtl crS,target		equivalent to		bcl 12,BI <sup>(4)</sup> ,target																						<b>bgtl</b>							
<b>bgtla</b>	bgtla crS,target		equivalent to		bcla 12,BI <sup>(4)</sup> ,target																						<b>bgtla</b>							
<b>bgtlr</b>	bgtlr crS,target		equivalent to		bclr 12,BI <sup>(4)</sup> ,target																						<b>bgtlr</b>							
<b>bgtlrl</b>	bgtlrl crS,target		equivalent to		bclrl 12,BI <sup>(4)</sup> ,target																						<b>bgtlrl</b>							
<b>bl</b>	18 (0x12)						LI																				0	1	I	<b>bl</b>				
<b>bla</b>	18 (0x12)						LI																				1	1	I	<b>bla</b>				
<b>ble</b>	ble crS,target		equivalent to		bc 4,BI <sup>(4)</sup> ,target																						<b>ble</b>							
<b>blea</b>	blea crS,target		equivalent to		bca 4,BI <sup>(4)</sup> ,target																						<b>blea</b>							
<b>blectr</b>	blectr crS,target		equivalent to		bcctr 4,BI <sup>(4)</sup> ,target																						<b>blectr</b>							
<b>blectrl</b>	blectrl crS,target		equivalent to		bcctrl 4,BI <sup>(4)</sup> ,target																						<b>blectrl</b>							
<b>blel</b>	blel crS,target		equivalent to		bcl 4,BI <sup>(4)</sup> ,target																						<b>blel</b>							
<b>blela</b>	blela crS,target		equivalent to		bcla 4,BI <sup>(4)</sup> ,target																						<b>blela</b>							
<b>blelr</b>	blelr crS,target		equivalent to		bclr 4,BI <sup>(4)</sup> ,target																						<b>blelr</b>							
<b>blelrl</b>	blelrl crS,target		equivalent to		bclrl 4,BI <sup>(4)</sup> ,target																						<b>blelrl</b>							
<b>blr</b>	blr <sup>(1)</sup>		equivalent to		bclr 20,0																						<b>blr</b>							
<b>blrl</b>	blrl <sup>(1)</sup>		equivalent to		bclrl 20,0																						<b>blrl</b>							
<b>blt</b>	blt crS,target		equivalent to		bc 12,BI,target																						<b>blt</b>							
<b>blta</b>	blta crS,target		equivalent to		bca 12,BI <sup>(3)</sup> ,target																						<b>blta</b>							



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>blctr</b>	blctr crS,target		equivalent to		bcctr 12,Bl <sup>(3)</sup> ,target																						<b>blctr</b>							
<b>blctr1</b>	blctr1 crS,target		equivalent to		bcctr1 12,Bl <sup>(3)</sup> ,target																						<b>blctr1</b>							
<b>bltl</b>	bltl crS,target		equivalent to		bcl 12,Bl <sup>(3)</sup> ,target																						<b>bltl</b>							
<b>bltla</b>	bltla crS,target		equivalent to		bcla 12,Bl <sup>(3)</sup> ,target																						<b>bltla</b>							
<b>bltlr</b>	bltlr crS,target		equivalent to		bclr 12,Bl <sup>(3)</sup> ,target																						<b>bltlr</b>							
<b>bltlrl</b>	bltlrl crS,target		equivalent to		bclrl 12,Bl <sup>(3)</sup> ,target																						<b>bltlrl</b>							
<b>bne</b>	bne crS,target		equivalent to		bc 4,Bl <sup>(3)</sup> ,target																						<b>bne</b>							
<b>bnea</b>	bnea crS,target		equivalent to		bca 4,Bl <sup>(3)</sup> ,target																						<b>bnea</b>							
<b>bnctr</b>	bnctr crS,target		equivalent to		bcctr 4,Bl <sup>(3)</sup> ,target																						<b>bnctr</b>							
<b>bnctr1</b>	bnctr1 crS,target		equivalent to		bcctr1 4,Bl <sup>(3)</sup> ,target																						<b>bnctr1</b>							
<b>bnel</b>	bnel crS,target		equivalent to		bcl 4,Bl <sup>(3)</sup> ,target																						<b>bnel</b>							
<b>bnela</b>	bnela crS,target		equivalent to		bcla 4,Bl <sup>(3)</sup> ,target																						<b>bnela</b>							
<b>bnelr</b>	bnelr crS,target		equivalent to		bclr 4,Bl <sup>(3)</sup> ,target																						<b>bnelr</b>							
<b>bnelrl</b>	bnelrl crS,target		equivalent to		bclrl 4,Bl <sup>(3)</sup> ,target																						<b>bnelrl</b>							
<b>bng</b>	bng crS,target		equivalent to		bc 4,Bl <sup>(4)</sup> ,target																						<b>bng</b>							
<b>bnga</b>	bnga crS,target		equivalent to		bca 4,Bl <sup>(4)</sup> ,target																						<b>bnga</b>							
<b>bngctr</b>	bngctr crS,target		equivalent to		bcctr 4,Bl <sup>(4)</sup> ,target																						<b>bngctr</b>							
<b>bngctr1</b>	bngctr1 crS,target		equivalent to		bcctr1 4,Bl <sup>(4)</sup> ,target																						<b>bngctr1</b>							
<b>bngl</b>	bngl crS,target		equivalent to		bcl 4,Bl <sup>(4)</sup> ,target																						<b>bngl</b>							
<b>bngla</b>	bngla crS,target		equivalent to		bcla 4,Bl <sup>(4)</sup> ,target																						<b>bngla</b>							
<b>bnglr</b>	bnglr crS,target		equivalent to		bclr 4,Bl <sup>(4)</sup> ,target																						<b>bnglr</b>							
<b>bnglrl</b>	bnglrl crS,target		equivalent to		bclrl 4,Bl <sup>(4)</sup> ,target																						<b>bnglrl</b>							
<b>bni</b>	bni crS,target		equivalent to		bc 4,Bl <sup>(3)</sup> ,target																						<b>bni</b>							
<b>bnla</b>	bnla crS,target		equivalent to		bca 4,Bl <sup>(3)</sup> ,target																						<b>bnla</b>							
<b>bnlctr</b>	bnlctr crS,target		equivalent to		bcctr 4,Bl <sup>(3)</sup> ,target																						<b>bnlctr</b>							





**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>bnlctrl</b>	bnlctrl crS,target		equivalent to		bcctrl 4,BI <sup>(3)</sup> ,target																						<b>bnlctrl</b>							
<b>bnll</b>	bnll crS,target		equivalent to		bcl 4,BI <sup>(3)</sup> ,target																						<b>bnll</b>							
<b>bnlla</b>	bnlla crS,target		equivalent to		bcla 4,BI <sup>(3)</sup> ,target																						<b>bnlla</b>							
<b>bnllr</b>	bnllr crS,target		equivalent to		bcrlr 4,BI <sup>(3)</sup> ,target																						<b>bnllr</b>							
<b>bnllrl</b>	bnllrl crS,target		equivalent to		bcrlrl 4,BI <sup>(3)</sup> ,target																						<b>bnllrl</b>							
<b>bns</b>	bns crS,target		equivalent to		bc 4,BI <sup>(5)</sup> ,target																						<b>bns</b>							
<b>bnsa</b>	bnsa crS,target		equivalent to		bca 4,BI <sup>(5)</sup> ,target																						<b>bnsa</b>							
<b>bnsctr</b>	bnsctr crS,target		equivalent to		bcctr 4,BI <sup>(5)</sup> ,target																						<b>bnsctr</b>							
<b>bnsctrl</b>	bnsctrl crS,target		equivalent to		bcctrl 4,BI <sup>(5)</sup> ,target																						<b>bnsctrl</b>							
<b>bnsi</b>	bnsi crS,target		equivalent to		bcl 4,BI <sup>(5)</sup> ,target																						<b>bnsi</b>							
<b>bnsia</b>	bnsia crS,target		equivalent to		bcla 4,BI <sup>(5)</sup> ,target																						<b>bnsia</b>							
<b>bnslr</b>	bnslr crS,target		equivalent to		bcrlr 4,BI <sup>(5)</sup> ,target																						<b>bnslr</b>							
<b>bnslrl</b>	bnslrl crS,target		equivalent to		bcrlrl 4,BI <sup>(5)</sup> ,target																						<b>bnslrl</b>							
<b>bnu</b>	bnu crS,target		equivalent to		bc 4,BI <sup>(5)</sup> ,target																						<b>bnu</b>							
<b>bnua</b>	bnua crS,target		equivalent to		bca 4,BI <sup>(5)</sup> ,target																						<b>bnua</b>							
<b>bnuctr</b>	bnuctr crS,target		equivalent to		bcctr 4,BI <sup>(5)</sup> ,target																						<b>bnuctr</b>							
<b>bnuctrl</b>	bnuctrl crS,target		equivalent to		bcctrl 4,BI <sup>(5)</sup> ,target																						<b>bnuctrl</b>							
<b>bnul</b>	bnul crS,target		equivalent to		bcl 4,BI <sup>(5)</sup> ,target																						<b>bnul</b>							
<b>bnula</b>	bnula crS,target		equivalent to		bcla 4,BI <sup>(5)</sup> ,target																						<b>bnula</b>							
<b>bnulr</b>	bnulr crS,target		equivalent to		bcrlr 4,BI <sup>(5)</sup> ,target																						<b>bnulr</b>							
<b>bnulrl</b>	bnulrl crS,target		equivalent to		bcrlrl 4,BI <sup>(5)</sup> ,target																						<b>bnulrl</b>							
<b>brinc</b>	04		rD		rA		rB		0		1		0		0		0		0		0		1		1		1		1		EVX		<b>brinc</b>	
<b>bso</b>	bso crS,target		equivalent to		bc 12,BI <sup>(5)</sup> ,target																						<b>bso</b>							
<b>bsoa</b>	bsoa crS,target		equivalent to		bca 12,BI <sup>(5)</sup> ,target																						<b>bsoa</b>							
<b>bsoctr</b>	bsoctr crS,target		equivalent to		bcctr 12,BI <sup>(5)</sup> ,target																						<b>bsoctr</b>							



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic																	
<b>bsoctrl</b>	bsoctrl crS,target equivalent to										bcctrl 12,BI <sup>(5)</sup> ,target																												<b>bsoctrl</b>												
<b>bsol</b>	bsol crS,target equivalent to										bcl 12,BI <sup>(5)</sup> ,target																													<b>bsol</b>											
<b>bsola</b>	bsola crS,target equivalent to										bcla 12,BI <sup>(5)</sup> ,target																													<b>bsola</b>											
<b>bsolr</b>	bsolr crS,target equivalent to										bc1r 12,BI <sup>(5)</sup> ,target																													<b>bsolr</b>											
<b>bsolrl</b>	bsolrl crS,target equivalent to										bc1rl 12,BI <sup>(5)</sup> ,target																													<b>bsolrl</b>											
<b>bt</b>	bt BI,target equivalent to										bc 12,BI,target																													<b>bt</b>											
<b>bta</b>	bta BI,target equivalent to										bca 12,BI,target																														<b>bta</b>										
<b>btctr</b>	btctr BI equivalent to										bcctr 12,BI																														<b>btctr</b>										
<b>btctrl</b>	btctrl BI equivalent to										bcctrl 12,BI																															<b>btctrl</b>									
<b>btl</b>	btl BI,target equivalent to										bcl 12,BI,target																															<b>btl</b>									
<b>btla</b>	btla BI,target equivalent to										bcla 12,BI,target																															<b>btla</b>									
<b>btlr</b>	btlr BI equivalent to										bc1r 12,BI																															<b>btlr</b>									
<b>btlrl</b>	btlrl BI equivalent to										bc1rl 12,BI																																<b>btlrl</b>								
<b>bun</b>	bun crS,target equivalent to										bc 12,BI <sup>(5)</sup> ,target																															<b>bun</b>									
<b>buna</b>	buna crS,target equivalent to										bca 12,BI <sup>(5)</sup> ,target																															<b>buna</b>									
<b>bunctr</b>	bunctr crS,target equivalent to										bcctr 12,BI <sup>(5)</sup> ,target																															<b>bunctr</b>									
<b>bunctrl</b>	bunctrl crS,target equivalent to										bcctrl 12,BI <sup>(5)</sup> ,target																															<b>bunctrl</b>									
<b>bunl</b>	bunl crS,target equivalent to										bcl 12,BI <sup>(5)</sup> ,target																																<b>bunl</b>								
<b>bunla</b>	bunla crS,target equivalent to										bcla 12,BI <sup>(5)</sup> ,target																																<b>bunla</b>								
<b>bunlr</b>	bunlr crS,target equivalent to										bc1r 12,BI <sup>(5)</sup> ,target																																<b>bunlr</b>								
<b>bunlrl</b>	bunlrl crS,target equivalent to										bc1rl 12,BI <sup>(5)</sup> ,target																																<b>bunlrl</b>								
<b>clrlslwi</b>	clrlslwi rA,rS,b,n (n ∈ b ∈ 31)										equivalent to										rlwinm rA,rS,n,b – n,31 – n																													<b>clrlslwi</b>	
<b>clrlwi</b>	clrlwi rA,rS,n (n < 32)										equivalent to										rlwinm rA,rS,0,n,31																														<b>clrlwi</b>
<b>clrrwi</b>	clrrwi rA,rS,n (n < 32)										equivalent to										rlwinm rA,rS,0,0,31 – n																														<b>clrrwi</b>
<b>cmp</b>	31 (0x1F)					crfD					/ L					rA					rB					0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 /					X					<b>cmp</b>															



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
<b>cmpi</b>	11 (0x0B)						crfD	/	L	rA					SIMM											D	<b>cmpi</b>								
<b>cmpl</b>	31 (0x1F)					/	L	rA					rB					///	0	0	0	0	1	0	0	0	0	0	0	0	0	/	X	<b>cmpl</b>	
<b>cmpli</b>	10 (0x0A)						crfD	/	L	rA					UIMM											D	<b>cmpli</b>								
<b>cmplw</b>	cmplw crD,rA,rB					equivalent to					cmpl crD,0,rA,rB																	<b>cmplw</b>							
<b>cmplwi</b>	cmplwi crD,rA,UIMM					equivalent to					cmpli crD,0,rA,UIMM																	<b>cmplwi</b>							
<b>cmpw</b>	cmpw crD,rA,rB					equivalent to					cmp crD,0,rA,rB																	<b>cmpw</b>							
<b>cmpwi</b>	cmpwi crD,rA,SIMM					equivalent to					cmpi crD,0,rA,SIMM																	<b>cmpwi</b>							
<b>cntlzw</b>	31 (0x1F)						rS					rA					///					0	0	0	0	0	1	1	0	1	0	0	X	<b>cntlzw</b>	
<b>cntlzw.</b>	31 (0x1F)						rS					rA					///					0	0	0	0	0	1	1	0	1	0	1	X	<b>cntlzw.</b>	
<b>crand</b>	19 (0x13)						crbD					crbA					crbB					0	1	0	0	0	0	0	0	0	1	/	XL	<b>crand</b>	
<b>crandc</b>	19 (0x13)						crbD					crbA					crbB					0	0	1	0	0	0	0	0	0	1	/	XL	<b>crandc</b>	
<b>crclr</b>	crclr bx					equivalent to					crxor bx,bx,bx																	<b>crclr</b>							
<b>creqv</b>	19 (0x13)						crbD					crbA					crbB					0	1	0	0	1	0	0	0	0	0	1	/	XL	<b>creqv</b>
<b>crmmove</b>	crmmove bx,by					equivalent to					cror bx,by,by																	<b>crmmove</b>							
<b>crnand</b>	19 (0x13)						crbD					crbA					crbB					0	0	1	1	1	0	0	0	0	1	/	XL	<b>crnand</b>	
<b>crnor</b>	19 (0x13)						crbD					crbA					crbB					0	0	0	0	1	0	0	0	0	1	/	XL	<b>crnor</b>	
<b>crnot</b>	crnot bx,by					equivalent to					crnor bx,by,by																	<b>crnot</b>							
<b>cror</b>	19 (0x13)						crbD					crbA					crbB					0	1	1	1	0	0	0	0	0	1	/	XL	<b>cror</b>	
<b>crorc</b>	19 (0x13)						crbD					crbA					crbB					0	1	1	0	1	0	0	0	0	1	/	XL	<b>crorc</b>	
<b>crset</b>	crset bx					equivalent to					creqv bx,bx,bx																	<b>crset</b>							
<b>crxor</b>	19 (0x13)						crbD					crbA					crbB					0	0	1	1	0	0	0	0	0	1	/	XL	<b>crxor</b>	
<b>dcba</b> <sup>(6)</sup>	31 (0x1F)						///					rA					rB					1	0	1	1	1	1	0	1	1	0	/	X	<b>dcba</b>	
<b>dcbf</b>	31 (0x1F)						///					rA					rB					0	0	0	1	0	1	0	1	1	0	/	X	<b>dcbf</b>	
<b>dcbi</b> <sup>(7)</sup>	31 (0x1F)						///					rA					rB					0	1	1	1	0	1	0	1	1	0	/	X	<b>dcbi</b>	
<b>dcblc</b>	31 (0x1F)						CT					rA					rB					0	1	1	0	0	0	0	1	1	0	0	X	<b>dcblc</b>	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>dcbst</b>	31 (0x1F)						///					rA					rB					0	0	0	0	1	1	0	1	1	0	/	X	<b>dcbst</b>
<b>dcbt</b>	31 (0x1F)						CT					rA					rB					0	1	0	0	0	1	0	1	1	0	/	X	<b>dcbt</b>
<b>dcbtls</b>	31 (0x1F)						CT					rA					rB					0	0	1	0	1	0	0	1	1	0	0	X	<b>dcbtls</b>
<b>dcbtst</b>	31 (0x1F)						CT					rA					rB					0	0	1	1	1	1	0	1	1	0	/	X	<b>dcbtst</b>
<b>dcbtstls</b>	31 (0x1F)						CT					rA					rB					0	0	1	0	0	0	0	1	1	0	0	X	<b>dcbtstls</b>
<b>dcbz</b>	31 (0x1F)						///					rA					rB					1	1	1	1	1	1	0	1	1	0	/	X	<b>dcbz</b>
<b>divw</b>	31 (0x1F)						rD					rA					rB					0	1	1	1	1	0	1	0	1	1	0	X	<b>divw</b>
<b>divw.</b>	31 (0x1F)						rD					rA					rB					0	1	1	1	1	0	1	0	1	1	1	X	<b>divw.</b>
<b>divwo</b>	31 (0x1F)						rD					rA					rB					1	1	1	1	1	0	1	0	1	1	0	X	<b>divwo</b>
<b>divwo.</b>	31 (0x1F)						rD					rA					rB					1	1	1	1	1	0	1	0	1	1	1	X	<b>divwo.</b>
<b>divwu</b>	31 (0x1F)						rD					rA					rB					0	1	1	1	0	0	1	0	1	1	0	X	<b>divwu</b>
<b>divwu.</b>	31 (0x1F)						rD					rA					rB					0	1	1	1	0	0	1	0	1	1	1	X	<b>divwu.</b>
<b>divwuo</b>	31 (0x1F)						rD					rA					rB					1	1	1	1	0	0	1	0	1	1	0	X	<b>divwuo</b>
<b>divwuo.</b>	31 (0x1F)						rD					rA					rB					1	1	1	1	0	0	1	0	1	1	1	X	<b>divwuo.</b>
<b>dss</b>	dss STRM						equivalent to										dss STRM,0																<b>dss</b>	
<b>efdabs</b>	04						rD					rA					///					0	1	0	1	1	1	0	0	1	0	0	EFX	<b>efdabs</b>
<b>efdadd</b>	04						rD					rA					rB					0	1	0	1	1	1	0	0	0	0	0	EFX	<b>efdadd</b>
<b>efdcfs</b>	04						rD					0	0	0	0	0	rB					0	1	0	1	1	1	0	1	1	1	1	EFX	<b>efdcfs</b>
<b>efdcfsf</b>	04						rD					///					rB					0	1	0	1	1	1	1	0	0	1	1	EFX	<b>efdcfsf</b>
<b>efdcfsi</b>	04						rD					///					rB					0	1	0	1	1	1	1	0	0	0	1	EFX	<b>efdcfsi</b>
<b>efdcfuf</b>	04						rD					///					rB					0	1	0	1	1	1	1	0	0	1	0	EFX	<b>efdcfuf</b>
<b>efdcfui</b>	04						rD					///					rB					0	1	0	1	1	1	1	0	0	0	0	EFX	<b>efdcfui</b>
<b>efdcmpcq</b>	04						crfD		/	/	rA					rB					0	1	0	1	1	1	0	1	1	1	0	EFX	<b>efdcmpcq</b>	
<b>efdcmpgt</b>	04						crfD		/	/	rA					rB					0	1	0	1	1	1	0	1	1	0	0	EFX	<b>efdcmpgt</b>	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>efdcmlpt</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	1	0	1	1	0	1	EFX	<b>efdcmlpt</b>	
<b>efdctsf</b>	04						rD					///							rB		0	1	0	1	1	1	1	0	1	1	1	EFX	<b>efdctsf</b>	
<b>efdctsi</b>	04						rD					///							rB		0	1	0	1	1	1	1	0	1	0	1	EFX	<b>efdctsi</b>	
<b>efdctsiz</b>	04						rD					///							rB		0	1	0	1	1	1	1	1	0	1	0	EFX	<b>efdctsiz</b>	
<b>efdctuf</b>	04						rD					///							rB		0	1	0	1	1	1	1	0	1	1	0	EFX	<b>efdctuf</b>	
<b>efdctui</b>	04						rD					///							rB		0	1	0	1	1	1	1	0	1	0	0	EFX	<b>efdctui</b>	
<b>efdctuiz</b>	04						rD					///							rB		0	1	0	1	1	1	1	1	0	0	0	EFX	<b>efdctuiz</b>	
<b>efddiv</b>	04						rD					rA							rB		0	1	0	1	1	1	0	1	0	0	1	EFX	<b>efddiv</b>	
<b>efdmul</b>	04						rD					rA							rB		0	1	0	1	1	1	0	1	0	0	0	EFX	<b>efdmul</b>	
<b>efdnabs</b>	04						rD					rA							///		0	1	0	1	1	1	0	0	1	0	1	EFX	<b>efdnabs</b>	
<b>efdneg</b>	04						rD					rA							///		0	1	0	1	1	1	0	0	1	1	0	EFX	<b>efdneg</b>	
<b>efdsbub</b>	04						rD					rA							rB		0	1	0	1	1	1	0	0	0	0	1	EFX	<b>efdsbub</b>	
<b>efdtsteg</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	1	1	1	1	1	0	EFX	<b>efdtsteg</b>	
<b>efdtstgt</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	1	1	1	1	1	0	0	EFX	<b>efdtstgt</b>
<b>efdtstlt</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	1	1	1	1	1	0	1	EFX	<b>efdtstlt</b>
<b>efsabs</b>	04						rD					rA							///		0	1	0	1	1	0	0	0	1	0	0	EFX	<b>efsabs</b>	
<b>efsadd</b>	04						rD					rA							rB		0	1	0	1	1	0	0	0	0	0	0	0	EFX	<b>efsadd</b>
<b>efscfd</b>	04						rD				0	0	0	0	0				rB		0	1	0	1	1	0	0	1	1	1	1	EFX	<b>efscfd</b>	
<b>efscfsf</b>	04						rD					///							rB		0	1	0	1	1	0	1	0	0	1	1	EFX	<b>efscfsf</b>	
<b>efscfsi</b>	04						rD					///							rB		0	1	0	1	1	0	1	0	0	0	1	EFX	<b>efscfsi</b>	
<b>efscfuf</b>	04						rD					///							rB		0	1	0	1	1	0	1	0	0	1	0	EFX	<b>efscfuf</b>	
<b>efscfui</b>	04						rD					///							rB		0	1	0	1	1	0	1	0	0	0	0	EFX	<b>efscfui</b>	
<b>efscmpeq</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	0	0	1	1	1	0	EFX	<b>efscmpeq</b>	
<b>efscmpgt</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	0	0	1	1	0	0	EFX	<b>efscmpgt</b>	
<b>efscmplt</b>	04						crfD	/	/			rA							rB		0	1	0	1	1	0	0	1	1	0	1	EFX	<b>efscmplt</b>	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>efsctsf</b>	04						rD			///						rB			0	1	0	1	1	0	1	0	1	1	1	1	1	EFX	<b>efsctsf</b>	
<b>efsctsi</b>	04						rD			///						rB			0	1	0	1	1	0	1	0	1	0	1	0	1	EFX	<b>efsctsi</b>	
<b>efsctsiz</b>	04						rD			///						rB			0	1	0	1	1	0	1	1	0	1	0	1	0	EFX	<b>efsctsiz</b>	
<b>efsctuf</b>	04						rD			///						rB			0	1	0	1	1	0	1	0	1	0	1	1	0	EFX	<b>efsctuf</b>	
<b>efsctui</b>	04						rD			///						rB			0	1	0	1	1	0	1	0	1	0	0	0	EFX	<b>efsctui</b>		
<b>efsctuiZ</b>	04						rD			///						rB			0	1	0	1	1	0	1	1	0	0	0	0	0	EFX	<b>efsctuiZ</b>	
<b>efsdiv</b>	04						rD			rA			rB			0	1	0	1	1	0	0	1	0	0	1	0	0	1	EFX	<b>efsdiv</b>			
<b>efsmul</b>	04						rD			rA			rB			0	1	0	1	1	0	0	1	0	0	0	0	0	0	EFX	<b>efsmul</b>			
<b>efsnabs</b>	04						rD			rA			///			0	1	0	1	1	0	0	0	0	1	0	1	EFX	<b>efsnabs</b>					
<b>efsneg</b>	04						rD			rA			///			0	1	0	1	1	0	0	0	0	1	1	0	EFX	<b>efsneg</b>					
<b>efssub</b>	04						rD			rA			rB			0	1	0	1	1	0	0	0	0	0	0	0	1	EFX	<b>efssub</b>				
<b>efststeq</b>	04						crfD		/	/	rA			rB			0	1	0	1	1	0	1	1	1	1	0	EFX	<b>efststeq</b>					
<b>efststgt</b>	04						crfD		/	/	rA			rB			0	1	0	1	1	0	1	1	1	0	0	EFX	<b>efststgt</b>					
<b>efststlt</b>	04						crfD		/	/	rA			rB			0	1	0	1	1	0	1	1	1	0	1	EFX	<b>efststlt</b>					
<b>eqv</b>	31 (0x1F)						rD			rA			rB			0	1	0	0	0	1	1	1	0	0	0	X	<b>eqv</b>						
<b>eqv.</b>	31 (0x1F)						rD			rA			rB			0	1	0	0	0	1	1	1	0	0	1	X	<b>eqv.</b>						
<b>evabs</b>	31 (0x1F)						rD			rA			///			0	1	0	0	0	0	0	1	0	0	0	EVX	<b>evabs</b>						
<b>evaddiw</b>	31 (0x1F)						rD			UIMM			rB			0	1	0	0	0	0	0	0	0	1	0	EVX	<b>evaddiw</b>						
<b>evaddsmiaaw</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	0	1	0	0	1	EVX	<b>evaddsmiaaw</b>						
<b>evaddssiaw</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	0	0	0	0	1	EVX	<b>evaddssiaw</b>						
<b>evaddumiaaw</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	0	1	0	0	0	EVX	<b>evaddumiaaw</b>						
<b>evaddusiaw</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	EVX	<b>evaddusiaw</b>					



**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic					
<b>evaddw</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	EVX	<b>evaddw</b>
<b>evand</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	0	0	0	0	1	0	0	0	0	1	EVX	<b>evand</b>	
<b>evandc</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	0	0	0	1	0	0	1	0	0	EVX	<b>evandc</b>		
<b>evcmpeq</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	0	0	0	1	1	0	1	0	0	EVX	<b>evcmpeq</b>		
<b>evcmpgts</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	0	0	0	1	1	0	0	0	1	EVX	<b>evcmpgts</b>		
<b>evcmpgtu</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	0	0	0	1	1	0	0	0	0	EVX	<b>evcmpgtu</b>		
<b>evcmplt</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	0	0	0	1	1	0	0	1	1	EVX	<b>evcmplt</b>		
<b>evcmpltu</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	0	0	0	1	1	0	0	1	0	EVX	<b>evcmpltu</b>		
<b>evcntlsw</b>	31 (0x1F)						rD			rA						///						0	1	0	0	0	0	0	0	0	1	1	1	0	EVX	<b>evcntlsw</b>			
<b>evcntlzw</b>	31 (0x1F)						rD			rA						///						0	1	0	0	0	0	0	0	0	1	1	0	1	EVX	<b>evcntlzw</b>			
<b>evdivws</b>	31 (0x1F)						rD			rA						rB						1	0	0	1	1	0	0	0	0	1	1	0	EVX	<b>evdivws</b>				
<b>evdivwu</b>	31 (0x1F)						rD			rA						rB						1	0	0	1	1	0	0	0	0	1	1	1	EVX	<b>evdivwu</b>				
<b>eveqv</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	0	0	1	1	0	0	1	EVX	<b>eveqv</b>				
<b>evextsb</b>	31 (0x1F)						rD			rA						///						0	1	0	0	0	0	0	0	1	0	1	0	EVX	<b>evextsb</b>				
<b>evextsh</b>	31 (0x1F)						rD			rA						///						0	1	0	0	0	0	0	0	1	0	1	1	EVX	<b>evextsh</b>				
<b>evfsabs</b>	31 (0x1F)						rD			rA						///						0	1	0	1	0	0	0	0	0	1	0	0	EVX	<b>evfsabs</b>				
<b>evfsadd</b>	31 (0x1F)						rD			rA						rB						0	1	0	1	0	0	0	0	0	0	0	0	0	EVX	<b>evfsadd</b>			
<b>evfscfsf</b>	31 (0x1F)						rD			///						rB						0	1	0	1	0	0	1	0	0	1	1	EVX	<b>evfscfsf</b>					
<b>evfscfsi</b>	31 (0x1F)						rD			///						rB						0	1	0	1	0	0	1	0	0	0	1	EVX	<b>evfscfsi</b>					
<b>evfscfuf</b>	31 (0x1F)						rD			///						rB						0	1	0	1	0	0	1	0	0	1	0	EVX	<b>evfscfuf</b>					
<b>evfscfui</b>	31 (0x1F)						rD			///						rB						0	1	0	1	0	0	1	0	0	0	0	EVX	<b>evfscfui</b>					
<b>evfscmpeq</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	1	0	0	0	1	1	1	0	EVX	<b>evfscmpeq</b>			
<b>evfscmpgt</b>	31 (0x1F)						crfD			/	/	rA						rB						0	1	0	1	0	0	0	1	1	0	0	EVX	<b>evfscmpgt</b>			



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>evfscmplt</b>	31 (0x1F)						crfD		/	/		rA					rB					0	1	0	1	0	0	0	1	1	0	1	EVX	<b>evfscmplt</b>
<b>evfsctsf</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	0	1	1	1	EVX	<b>evfsctsf</b>	
<b>evfsctsi</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	0	1	0	1	EVX	<b>evfsctsi</b>	
<b>evfsctsiz</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	1	0	1	0	EVX	<b>evfsctsiz</b>	
<b>evfsctuf</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	0	1	1	0	EVX	<b>evfsctuf</b>	
<b>evfsctui</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	0	1	0	0	EVX	<b>evfsctui</b>	
<b>evfsctuiZ</b>	31 (0x1F)						rD				///					rB					0	1	0	1	0	0	1	1	0	0	0	EVX	<b>evfsctuiZ</b>	
<b>evfsdiv</b>	31 (0x1F)						rD				rA					rB					0	1	0	1	0	0	0	1	0	0	1	EVX	<b>evfsdiv</b>	
<b>evfsmul</b>	31 (0x1F)						rD				rA					rB					0	1	0	1	0	0	0	1	0	0	0	EVX	<b>evfsmul</b>	
<b>evfsnabs</b>	31 (0x1F)						rD				rA					///					0	1	0	1	0	0	0	0	1	0	1	EVX	<b>evfsnabs</b>	
<b>evfsneg</b>	31 (0x1F)						rD				rA					///					0	1	0	1	0	0	0	0	1	1	0	EVX	<b>evfsneg</b>	
<b>evfssub</b>	31 (0x1F)						rD				rA					rB					0	1	0	1	0	0	0	0	0	0	1	EVX	<b>evfssub</b>	
<b>evfststeq</b>	31 (0x1F)						crfD		/	/		rA					rB					0	1	0	1	0	0	1	1	1	1	0	EVX	<b>evfststeq</b>
<b>evfststgt</b>	31 (0x1F)						crfD		/	/		rA					rB					0	1	0	1	0	0	1	1	1	0	0	EVX	<b>evfststgt</b>
<b>evfststlt</b>	31 (0x1F)						crfD		/	/		rA					rB					0	1	0	1	0	0	1	1	1	0	1	EVX	<b>evfststlt</b>
<b>evlidd</b>	31 (0x1F)						rD				rA					UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	0	0	0	1	EVX	<b>evlidd</b>
<b>evliddx</b>	31 (0x1F)						rD				rA					rB					0	1	1	0	0	0	0	0	0	0	0	0	EVX	<b>evliddx</b>
<b>evldh</b>	31 (0x1F)						rD				rA					UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	1	0	1	EVX	<b>evldh</b>	
<b>evldhx</b>	31 (0x1F)						rD				rA					rB					0	1	1	0	0	0	0	0	1	0	0	EVX	<b>evldhx</b>	
<b>evldw</b>	31 (0x1F)						rD				rA					UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	0	1	1	EVX	<b>evldw</b>	
<b>evldwx</b>	31 (0x1F)						rD				rA					rB					0	1	1	0	0	0	0	0	0	1	0	EVX	<b>evldwx</b>	
<b>evlhhesplat</b>	31 (0x1F)						rD				rA					UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	1	0	0	1	EVX	<b>evlhhesplat</b>	
<b>evlhhesplat<sub>x</sub></b>	31 (0x1F)						rD				rA					rB					0	1	1	0	0	0	0	1	0	0	0	EVX	<b>evlhhesplat<sub>x</sub></b>	





Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evlhossplat	31 (0x1F)					rD					rA					UIMM <sup>(9)</sup>					0	1	1	0	0	0	0	1	1	1	1	EVX	evlhossplat	
evlhossplatx	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	0	1	1	1	0	EVX	evlhossplatx	
evlhousplat	31 (0x1F)					rD					rA					UIMM <sup>(9)</sup>					0	1	1	0	0	0	0	1	1	0	1	EVX	evlhousplat	
evlhousplatx	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	0	1	1	0	0	EVX	evlhousplatx	
evlwhe	31 (0x1F)					rD					rA					UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	0	0	0	1	EVX	evlwhe	
evlwhex	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	1	0	0	0	0	EVX	evlwhex	
evlw hos	31 (0x1F)					rD					rA					UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	0	1	1	1	EVX	evlw hos	
evlw hosx	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	1	0	1	1	0	EVX	evlw hosx	
evlw hou	31 (0x1F)					rD					rA					UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	0	1	0	1	EVX	evlw hou	
evlw hou x	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	1	0	1	0	0	EVX	evlw hou x	
evlw hsplat	31 (0x1F)					rD					rA					UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	1	1	0	1	EVX	evlw hsplat	
evlw hsplat x	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	1	1	1	0	0	EVX	evlw hsplat x	
evlw wsplat	31 (0x1F)					rD					rA					UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	1	0	0	1	EVX	evlw wsplat	
evlw wsplat x	31 (0x1F)					rD					rA					rB					0	1	1	0	0	0	1	1	0	0	0	EVX	evlw wsplat x	
evmergehi	31 (0x1F)					rD					rA					rB					0	1	0	0	0	1	0	1	1	0	0	EVX	evmergehi	
evmergehilo	31 (0x1F)					rD					rA					rB					0	1	0	0	0	1	0	1	1	1	0	EVX	evmergehilo	
evmergeo	31 (0x1F)					rD					rA					rB					0	1	0	0	0	1	0	1	1	0	1	EVX	evmergeo	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic					
evmgerloh i	31 (0x1F)						rD						rA						rB						0	1	0	0	0	1	0	1	1	1	1	1	EVX	evmgerloh i	
evmhgsmf aa	31 (0x1F)						rD						rA						rB						1	0	1	0	0	1	0	1	0	1	1	1	EVX	evmhgsmf aa	
evmhgsmf an	31 (0x1F)						rD						rA						rB						1	0	1	1	0	1	0	1	0	1	1	1	EVX	evmhgsmf an	
evmhgsmi aa	31 (0x1F)						rD						rA						rB						1	0	1	0	0	1	0	1	0	0	1	1	EVX	evmhgsmi aa	
evmhgsmi an	31 (0x1F)						rD						rA						rB						1	0	1	1	0	1	0	1	0	0	1	1	EVX	evmhgsmi an	
evmhgumi aa	31 (0x1F)						rD						rA						rB						1	0	1	0	0	1	0	1	0	0	0	1	1	EVX	evmhgumi aa
evmhgumi an	31 (0x1F)						rD						rA						rB						1	0	1	1	0	1	0	1	0	0	0	1	1	EVX	evmhgumi an
evmhesmf	31 (0x1F)						rD						rA						rB						1	0	0	0	0	0	0	0	1	0	1	1	EVX	evmhesmf	
evmhesmfa	31 (0x1F)						rD						rA						rB						1	0	0	0	0	1	0	1	0	1	1	EVX	evmhesmfa		
evmhesmfa aw	31 (0x1F)						rD						rA						rB						1	0	1	0	0	0	0	1	0	1	1	EVX	evmhesmfa aw		
evmhesmfa nw	31 (0x1F)						rD						rA						rB						1	0	1	1	0	0	0	1	0	1	1	EVX	evmhesmfa nw		
evmhesmi	31 (0x1F)						rD						rA						rB						1	0	0	0	0	0	0	1	0	0	1	EVX	evmhesmi		
evmhesmia	31 (0x1F)						rD						rA						rB						1	0	0	0	0	1	0	1	0	0	1	EVX	evmhesmia		
evmhesmia aw	31 (0x1F)						rD						rA						rB						1	0	1	0	0	0	0	1	0	0	1	EVX	evmhesmia aw		
evmhesmia nw	31 (0x1F)						rD						rA						rB						1	0	1	1	0	0	0	1	0	0	1	EVX	evmhesmia nw		



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic		
evmhessf	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	0	0	0	0	0	0	1	1	EVX	evmhessf
evmhessfa	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	1	0	0	0	0	0	1	1	EVX	evmhessfa
evmhessfaaw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	0	0	0	0	0	1	1	EVX	evmhessfaaw
evmhessfanw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	0	0	0	0	0	1	1	EVX	evmhessfanw
evmhessiaaw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	0	0	0	0	0	0	1	EVX	evmhessiaaw
evmhessianw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	0	0	0	0	0	0	1	EVX	evmhessianw
evmheumi	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	0	0	0	0	1	0	0	0	EVX	evmheumi
evmheumia	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	1	0	1	0	0	0	0	EVX	evmheumia	
evmheumiaaw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	0	1	0	0	0	0	EVX	evmheumiaaw	
evmheumianw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	1	0	0	0	0	0	EVX	evmheumianw	
evmheusiaaw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evmheusiaaw
evmheusianw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	0	0	0	0	0	0	0	EVX	evmheusianw
evmhogsmfaa	31 (0x1F)					rD					rA					rB					1	0	1	0	0	1	0	1	1	1	1	1	1	EVX	evmhogsmfaa	
evmhogsmfan	31 (0x1F)					rD					rA					rB					1	0	1	1	0	1	0	1	1	1	1	1	1	EVX	evmhogsmfan	
evmhogsmiaa	31 (0x1F)					rD					rA					rB					1	0	1	0	0	1	0	1	1	1	0	1	EVX	evmhogsmiaa		



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhogsmi an	31 (0x1F)					rD					rA					rB					1	0	1	1	0	1	0	1	1	0	1	EVX	evmhogs mian	
evmhogumi aa	31 (0x1F)					rD					rA					rB					1	0	1	0	0	1	0	1	1	0	0	EVX	evmhogu miaa	
evmhogumi an	31 (0x1F)					rD					rA					rB					1	0	1	1	0	1	0	1	1	0	0	EVX	evmhogu mian	
evmhosmf	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	0	0	1	1	1	1	EVX	evmhosm f
evmhosmfa	31 (0x1F)					rD					rA					rB					1	0	0	0	0	1	0	1	1	1	1	EVX	evmhosm fa	
evmhosmfa aw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	1	1	1	1	EVX	evmhosm faaw	
evmhosmfa nw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	1	1	1	1	EVX	evmhosm fanw	
evmhosmi	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	0	1	1	0	1	EVX	evmhosm i	
evmhosmia	31 (0x1F)					rD					rA					rB					1	0	0	0	0	1	0	1	1	0	1	EVX	evmhosm ia	
evmhosmia aw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	1	1	0	1	EVX	evmhosm iaaw	
evmhosmia nw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	1	1	0	1	EVX	evmhosm ianw	
evmhossf	31 (0x1F)					rD					rA					rB					1	0	0	0	0	0	0	0	1	1	1	EVX	evmhossf	
evmhossfa	31 (0x1F)					rD					rA					rB					1	0	0	0	0	1	0	0	1	1	1	EVX	evmhossf a	
evmhossfa aw	31 (0x1F)					rD					rA					rB					1	0	1	0	0	0	0	0	1	1	1	EVX	evmhossf aaw	
evmhossfa nw	31 (0x1F)					rD					rA					rB					1	0	1	1	0	0	0	0	1	1	1	EVX	evmhossf anw	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic			
<b>evmhossiaaw</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	1	0	1	EVX	<b>evmhossiaaw</b>					
<b>evmhossianw</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	1	0	1	EVX	<b>evmhossianw</b>						
<b>evmhoumi</b>	31 (0x1F)						rD			rA			rB			1	0	0	0	0	0	0	0	0	0	1	1	0	0	EVX	<b>evmhoumi</b>						
<b>evmhoumia</b>	31 (0x1F)						rD			rA			rB			1	0	0	0	0	0	1	0	1	1	0	0	EVX	<b>evmhoumia</b>								
<b>evmhoumiaaw</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	0	0	0	0	1	1	0	0	EVX	<b>evmhoumiaaw</b>								
<b>evmhoumianw</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	0	0	0	1	1	0	0	EVX	<b>evmhoumianw</b>									
<b>evmhousiaaw</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	0	0	0	0	0	1	0	0	EVX	<b>evmhousiaaw</b>								
<b>evmhousianw</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	0	0	0	0	0	1	0	0	EVX	<b>evmhousianw</b>								
<b>evmr</b>	evmr rD,rA equivalent to						evor rD,rA,rA																														<b>evmr</b>
<b>evmra</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	0	0	0	1	0	0	EVX	<b>evmra</b>								
<b>evmwhgsmfaa</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	1	1	0	1	1	1	1	1	EVX	<b>evmwhgsmfaa</b>								
<b>evmwhgsmfan</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	1	0	1	1	1	1	1	EVX	<b>evmwhgsmfan</b>									
<b>evmwhgsmiaa</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	1	1	0	1	1	0	1	EVX	<b>evmwhgsmiaa</b>									
<b>evmwhgsmian</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	EVX	<b>evmwhgsmian</b>									
<b>evmwhgssfaa</b>	31 (0x1F)						rD			rA			rB			1	0	1	0	1	1	0	0	1	1	1	EVX	<b>evmwhgssfaa</b>									
<b>evmwhgssfan</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	1	0	1	0	1	1	1	EVX	<b>evmwhgssfan</b>									



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic		
evmwhgum iaa	31 (0x1F)							rD						rA											rB										EVX	evmwhgu miaa
evmwhgum ian	31 (0x1F)							rD						rA											rB										EVX	evmwhgu mian
evmwhsmf	31 (0x1F)							rD						rA											rB										EVX	evmwhs mf
evmwhsmf a	31 (0x1F)							rD						rA											rB										EVX	evmwhs mfa
evmwhsmf aaw	31 (0x1F)							rD						rA											rB										EVX	evmwhs mfaaw
evmwhsmf anw	31 (0x1F)							rD						rA											rB										EVX	evmwhs mfanw
evmwhsmi	31 (0x1F)							rD						rA											rB										EVX	evmwhs mi
evmwhsmi a	31 (0x1F)							rD						rA											rB										EVX	evmwhs mia
evmwhsmi aaw	31 (0x1F)							rD						rA											rB										EVX	evmwhs miaaw
evmwhsmi anw	31 (0x1F)							rD						rA											rB										EVX	evmwhs mianw
evmwhssf	31 (0x1F)							rD						rA											rB										EVX	evmwhss f
evmwhssfa	31 (0x1F)							rD						rA											rB										EVX	evmwhss fa
evmwhssfa aw	31 (0x1F)							rD						rA											rB										EVX	evmwhss faaw
evmwhssfa nw	31 (0x1F)							rD						rA											rB										EVX	evmwhss fanw
evmwhssia nw	31 (0x1F)							rD						rA											rB										EVX	evmwhss ianw



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmwhssm aaw	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	0	0	1	0	1	EVX	evmwhss maaw	
evmwhumi	31 (0x1F)					rD					rA					rB					1	0	0	0	1	0	0	1	1	0	0	EVX	evmwhu mi	
evmwhumi a	31 (0x1F)					rD					rA					rB					1	0	0	0	1	1	0	1	1	0	0	EVX	evmwhu mia	
evmwhusia aw	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	0	0	1	0	0	EVX	evmwhus iaaw	
evmwhusia nw	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	0	0	1	0	0	EVX	evmwhus ianw	
evmwlsmf	31 (0x1F)					rD					rA					rB					1	0	0	0	1	0	0	1	0	1	1	EVX	evmwls mf	
evmwlsmf a	31 (0x1F)					rD					rA					rB					1	0	0	0	1	1	0	1	0	1	1	EVX	evmwls fa	
evmwlsmf aw	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	0	1	0	1	1	EVX	evmwls faaw	
evmwlsmf nw	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	0	1	0	1	1	EVX	evmwls fanw	
evmwlsmia aw	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	0	1	0	0	1	EVX	evmwls iaaw	
evmwlsmia nw	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	0	1	0	0	1	EVX	evmwls ianw	
evmwlsf	31 (0x1F)					rD					rA					rB					1	0	0	0	1	0	0	0	0	1	1	EVX	evmwls sf	
evmwlsf a	31 (0x1F)					rD					rA					rB					1	0	0	0	1	1	0	0	0	1	1	EVX	evmwls sfa	
evmwlsf aaw	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	0	0	0	1	1	EVX	evmwls sfaaw	
evmwlsf anw	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	0	0	0	1	1	EVX	evmwls sfanw	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
<b>evmwlsiaaw</b>	31 (0x1F)						rD			rA						rB			1	0	1	0	1	0	0	0	0	0	0	0	0	1	EVX	<b>evmwlsiaaw</b>	
<b>evmwlsianw</b>	31 (0x1F)						rD			rA						rB			1	0	1	1	1	1	0	0	0	0	0	0	0	1	EVX	<b>evmwlsianw</b>	
<b>evmwлумi</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	0	0	1	0	0	0	0	0	0	EVX	<b>evmwлумi</b>	
<b>evmwлумia</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	1	0	1	0	0	0	0	0	0	EVX	<b>evmwлумia</b>	
<b>evmwлумiaaw</b>	31 (0x1F)						rD			rA						rB			1	0	1	0	1	0	0	1	0	0	0	0	0	0	EVX	<b>evmwлумiaaw</b>	
<b>evmwлумianw</b>	31 (0x1F)						rD			rA						rB			1	0	1	1	1	0	0	1	0	0	0	0	0	0	EVX	<b>evmwлумianw</b>	
<b>evmwлusiaaw</b>	31 (0x1F)						rD			rA						rB			1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	EVX	<b>evmwлusiaaw</b>
<b>evmwлusianw</b>	31 (0x1F)						rD			rA						rB			1	0	1	1	1	0	0	0	0	0	0	0	0	0	0	EVX	<b>evmwлusianw</b>
<b>evmwsmf</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	0	1	1	0	1	1	1	1	EVX	<b>evmwsmf</b>		
<b>evmwsmfa</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	1	1	1	1	0	1	1	1	1	EVX	<b>evmwsmfa</b>	
<b>evmwsmfaa</b>	31 (0x1F)						rD			rA						rB			1	0	1	0	1	0	1	1	0	1	1	0	1	1	EVX	<b>evmwsmfaa</b>	
<b>evmwsmfan</b>	31 (0x1F)						rD			rA						rB			1	0	1	1	1	0	1	1	0	1	1	0	1	1	EVX	<b>evmwsmfan</b>	
<b>evmwsmi</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	0	1	1	0	1	1	0	0	1	EVX	<b>evmwsmi</b>	
<b>evmwsmia</b>	31 (0x1F)						rD			rA						rB			1	0	0	0	1	1	1	1	0	0	1	0	0	1	EVX	<b>evmwsmia</b>	
<b>evmwsmiaa</b>	31 (0x1F)						rD			rA						rB			1	0	1	0	1	0	1	1	0	0	0	0	1	EVX	<b>evmwsmiaa</b>		
<b>evmwsmian</b>	31 (0x1F)						rD			rA						rB			1	0	1	1	1	0	1	1	0	1	0	0	1	EVX	<b>evmwsmian</b>		





**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>evmwssf</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	1	0	1	0	0	1	1	EVX	<b>evmwssf</b>	
<b>evmwssfa</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	1	1	1	0	0	1	1	EVX	<b>evmwssfa</b>	
<b>evmwssfaa</b>	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	1	0	0	1	1	EVX	<b>evmwssfaa</b>	
<b>evmwssfan</b>	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	1	0	0	1	1	EVX	<b>evmwssfan</b>	
<b>evmwumi</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	1	0	1	1	0	0	0	EVX	<b>evmwumi</b>	
<b>evmwumia</b>	31 (0x1F)					rD					rA					rB					1	0	0	0	1	1	1	1	0	0	0	EVX	<b>evmwumia</b>	
<b>evmwumiaa</b>	31 (0x1F)					rD					rA					rB					1	0	1	0	1	0	1	1	0	0	0	EVX	<b>evmwumiaa</b>	
<b>evmwumian</b>	31 (0x1F)					rD					rA					rB					1	0	1	1	1	0	1	1	0	0	0	EVX	<b>evmwumian</b>	
<b>evnand</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	1	1	1	0	EVX	<b>evnand</b>	
<b>evneg</b>	31 (0x1F)					rD					rA					///					0	1	0	0	0	0	0	1	0	0	1	EVX	<b>evneg</b>	
<b>evnor</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	1	0	0	0	EVX	<b>evnor</b>	
<b>evnot</b>	evnot rD,rA					equivalent to					evnor rD,rA,rA																				<b>evnot</b>			
<b>evor</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	0	1	1	1	EVX	<b>evor</b>	
<b>evorc</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	1	1	0	1	1	EVX	<b>evorc</b>	
<b>evrlw</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	1	0	1	0	0	0	EVX	<b>evrlw</b>	
<b>evrlwi</b>	31 (0x1F)					rD					rA					UIMM					0	1	0	0	0	1	0	1	0	1	0	EVX	<b>evrlwi</b>	
<b>evrndw</b>	31 (0x1F)					rD					rA					UIMM					0	1	0	0	0	0	1	1	0	0	EVX	<b>evrndw</b>		
<b>evsel</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	1	1	1	1	crfS			EVX	<b>evsel</b>	
<b>evslw</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	1	0	0	1	0	0	EVX	<b>evslw</b>	
<b>evslwi</b>	31 (0x1F)					rD					rA					UIMM					0	1	0	0	0	1	0	0	1	1	0	EVX	<b>evslwi</b>	
<b>evsplatfi</b>	31 (0x1F)					rD					SIMM					///					0	1	0	0	0	1	0	1	0	1	1	EVX	<b>evsplatfi</b>	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>evsplati</b>	31 (0x1F)						rD			SIMM						///						0	1	0	0	0	1	0	1	0	0	1	EVX	<b>evsplati</b>
<b>evsrwis</b>	31 (0x1F)						rD			rA						UIMM						0	1	0	0	0	1	0	0	0	1	1	EVX	<b>evsrwis</b>
<b>evsrwiu</b>	31 (0x1F)						rD			rA						UIMM						0	1	0	0	0	1	0	0	0	1	0	EVX	<b>evsrwiu</b>
<b>evsrws</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	1	0	0	0	0	1	EVX	<b>evsrws</b>
<b>evsrwu</b>	31 (0x1F)						rD			rA						rB						0	1	0	0	0	1	0	0	0	0	0	EVX	<b>evsrwu</b>
<b>evstdd</b>	31 (0x1F)						rD			rA						UIMM <sup>(8)</sup>						0	1	1	0	0	1	0	0	0	0	1	EVX	<b>evstdd</b>
<b>evstddx</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	0	0	0	0	0	EVX	<b>evstddx</b>
<b>evstdh</b>	31 (0x1F)						rS			rA						UIMM <sup>(8)</sup>						0	1	1	0	0	1	0	0	1	0	1	EVX	<b>evstdh</b>
<b>evstdhx</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	0	0	1	0	0	EVX	<b>evstdhx</b>
<b>evstdw</b>	31 (0x1F)						rS			rA						UIMM <sup>(8)</sup>						0	1	1	0	0	1	0	0	0	1	1	EVX	<b>evstdw</b>
<b>evstdwx</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	0	0	0	1	0	EVX	<b>evstdwx</b>
<b>evstwhe</b>	31 (0x1F)						rS			rA						UIMM <sup>(10)</sup>						0	1	1	0	0	1	1	0	0	0	1	EVX	<b>evstwhe</b>
<b>evstwhex</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	1	0	0	0	0	EVX	<b>evstwhex</b>
<b>evstwho</b>	31 (0x1F)						rS			rA						UIMM <sup>(10)</sup>						0	1	1	0	0	1	1	0	1	0	1	EVX	<b>evstwho</b>
<b>evstwhox</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	1	0	1	0	0	EVX	<b>evstwhox</b>
<b>evstwwe</b>	31 (0x1F)						rS			rA						UIMM <sup>(10)</sup>						0	1	1	0	0	1	1	1	0	0	1	EVX	<b>evstwwe</b>
<b>evstwwex</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	1	1	0	0	0	EVX	<b>evstwwex</b>
<b>evstwwo</b>	31 (0x1F)						rS			rA						UIMM <sup>(10)</sup>						0	1	1	0	0	1	1	1	1	0	1	EVX	<b>evstwwo</b>
<b>evstwwox</b>	31 (0x1F)						rS			rA						rB						0	1	1	0	0	1	1	1	1	0	0	EVX	<b>evstwwox</b>
<b>evsubfsmiaaw</b>	31 (0x1F)						rD			rA						///						1	0	0	1	1	0	0	1	0	1	1	EVX	<b>evsubfsmiaaw</b>
<b>evsubfssiaaw</b>	31 (0x1F)						rD			rA						///						1	0	0	1	1	0	0	0	0	1	1	EVX	<b>evsubfssiaaw</b>
<b>evsubfumiaaw</b>	31 (0x1F)						rD			rA						///						1	0	0	1	1	0	0	1	0	1	0	EVX	<b>evsubfumiaaw</b>



**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic																
<b>evsubfusiaaw</b>	31 (0x1F)					rD					rA					///					1	0	0	1	1	0	0	0	0	0	1	0	EVX	<b>evsubfusiaw</b>																
<b>evsubfw</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	0	0	0	0	1	0	0	EVX	<b>evsubfw</b>															
<b>evsubifw</b>	31 (0x1F)					rD					UIMM					rB					0	1	0	0	0	0	0	0	0	1	1	0	EVX	<b>evsubifw</b>																
<b>evsubiw</b>	evsubiw rD,rB,UIMM										equivalent to										evsubifw rD,UIMM,rB																											<b>evsubiw</b>		
<b>evsubw</b>	evsubw rD,rB,rA										equivalent to										evsubfw rD,rA,rB																												<b>evsubw</b>	
<b>evxor</b>	31 (0x1F)					rD					rA					rB					0	1	0	0	0	0	0	1	0	1	1	0	EVX	<b>evxor</b>																
<b>extlwi</b>	extlwi rA,rS,n,b (n > 0)										equivalent to										rlwinm rA,rS,b,0,n - 1																											<b>extlwi</b>		
<b>extrwi</b>	extrwi rA,rS,n,b (n > 0)										equivalent to										rlwinm rA,rS,b + n, 32 - n,31																												<b>extrwi</b>	
<b>extsb</b>	31 (0x1F)					rS					rA					///					1	1	1	0	1	1	1	1	0	1	0	0	0	X	<b>extsb</b>															
<b>extsb.</b>	31 (0x1F)					rS					rA					///					1	1	1	0	1	1	1	1	0	1	0	1	X	<b>extsb.</b>																
<b>extsh</b>	31 (0x1F)					rS					rA					///					1	1	1	0	0	1	1	0	1	0	0	X	<b>extsh</b>																	
<b>extsh.</b>	31 (0x1F)					rS					rA					///					1	1	1	0	0	1	1	0	1	0	1	X	<b>extsh.</b>																	
<b>fres</b> <sup>(6)</sup>	59(0x3B)					frD					///					frB					///					1	1	0	0	0	0	A	<b>fres</b>																	
<b>fres.</b> <sup>(6)</sup>	59(0x3B)					frD					///					frB					///					1	1	0	0	0	1	A	<b>fres.</b>																	
<b>fsel</b> <sup>(6)</sup>	63(0x3F)					frD					frA					frB					frC					1	0	1	1	1	0	A	<b>fsel</b>																	
<b>fsel.</b> <sup>(6)</sup>	63(0x3F)					frD					frA					frB					frC					1	0	1	1	1	1	A	<b>fsel.</b>																	
<b>icbi</b>	31 (0x1F)					///					rA					rB					1	1	1	1	0	1	0	1	1	0	/	X	<b>icbi</b>																	
<b>icblc</b>	31 (0x1F)					CT					rA					rB					0	0	1	1	1	0	0	1	1	0	0	X	<b>icblc</b>																	
<b>icbt</b>	31 (0x1F)					CT					rA					rB					0	0	0	0	0	1	0	1	1	0	/	X	<b>icbt</b>																	
<b>icbtls</b>	31 (0x1F)					CT					rA					rB					0	1	1	1	1	0	0	1	1	0	0	X	<b>icbtls</b>																	
<b>inslwi</b>	inslwi rA,rS,n,b (n > 0)										equivalent to										rlwimi rA,rS,32 - b,b,(b + n) - 1																											<b>inslwi</b>		
<b>insrwi</b>	insrwi rA,rS,n,b (n > 0)										equivalent to										rlwimi rA,rS,32 - (b + n),b,(b + n) - 1																													<b>insrwi</b>
<b>isel</b>	31 (0x1F)					rD					rA					rB					crb					0	1	1	1	1	0	X	<b>isel</b>																	
<b>iseleq</b>	iseleq rD,rA,rB										equivalent to										isel rD,rA,rB,2																												<b>iseleq</b>	



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>iselgt</b>	iselgt rD,rA,rB equivalent to isel rD,rA,rB,1																															<b>iselgt</b>		
<b>isellt</b>	isellt rD,rA,rB equivalent to isel rD,rA,rB,0																															<b>isellt</b>		
<b>isync</b>	19 (0x13)					///											0	0	1	0	0	1	0	1	1	0	/	XL	<b>isync</b>					
<b>la</b>	la rD,d(rA) equivalent to addi rD,rA,d																															<b>la</b>		
<b>lbz</b>	34(0x22)					rD			rA			D											D	<b>lbz</b>										
<b>lbzu</b>	35(0x23)					rD			rA			D											D	<b>lbzu</b>										
<b>lbzux</b>	31 (0x1F)					rD			rA			rB			0	0	0	1	1	1	0	1	1	1	/	X	<b>lbzux</b>							
<b>lbzx</b>	31 (0x1F)					rD			rA			rB			0	0	0	1	0	1	0	1	1	1	/	X	<b>lbzx</b>							
<b>lha</b>	42(0x2A)					rD			rA			D											D	<b>lha</b>										
<b>lhau</b>	43(0x2B)					rD			rA			D											D	<b>lhau</b>										
<b>lhaux</b>	31 (0x1F)					rD			rA			rB			0	1	0	1	1	1	0	1	1	1	/	X	<b>lhaux</b>							
<b>lhax</b>	31 (0x1F)					rD			rA			rB			0	1	0	1	0	1	0	1	1	1	/	X	<b>lhax</b>							
<b>lhbrx</b>	31 (0x1F)					rD			rA			rB			1	1	0	0	0	1	0	1	1	0	/	X	<b>lhbrx</b>							
<b>lhz</b>	40(0x28)					rD			rA			D											D	<b>lhz</b>										
<b>lhzu</b>	41(0x29)					rD			rA			D											D	<b>lhzu</b>										
<b>lhzux</b>	31 (0x1F)					rD			rA			rB			0	1	0	0	1	1	0	1	1	1	/	X	<b>lhzux</b>							
<b>lhzx</b>	31 (0x1F)					rD			rA			rB			0	1	0	0	0	1	0	1	1	1	/	X	<b>lhzx</b>							
<b>li</b>	li rD,value equivalent to addi rD,0,value																															<b>li</b>		
<b>lis</b>	lis rD,value equivalent to addis rD,0,value																															<b>lis</b>		
<b>lmw</b>	46(0x2E)					rD			rA			D											D	<b>lmw</b>										
<b>lwarx</b>	31 (0x1F)					rD			rA			rB			0	0	0	0	0	1	0	1	0	0	/	X	<b>lwarx</b>							
<b>lwbrx</b>	31 (0x1F)					rD			rA			rB			1	0	0	0	0	1	0	1	1	0	/	X	<b>lwbrx</b>							
<b>lwz</b>	32 (0x20)					rD			rA			D											D	<b>lwz</b>										
<b>lwzu</b>	33 (0x21)					rD			rA			D											D	<b>lwzu</b>										
<b>lwzux</b>	31 (0x1F)					rD			rA			rB			0	0	0	0	1	1	0	1	1	1	/	X	<b>lwzux</b>							



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>lwzx</b>	31 (0x1F)					rD			rA					rB					0	0	0	0	0	1	0	1	1	1	/	X	<b>lwzx</b>			
<b>mbar</b>	31 (0x1F)					MO			///					1	1	0	1	0	1	0	1	1	0	/	X	<b>mbar</b>								
<b>mcrf</b>	19 (0x13)					crfD		//		crfS			///					0	0	0	0	0	0	0	0	0	0	0	0	0	0	/	XL	<b>mcrf</b>
<b>mcrxr</b>	31 (0x1F)					crfD		///					1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	/	X	<b>mcrxr</b>	
<b>mfcrr</b>	mfcrr rS					equivalent to					mfcrr 0xFF,rS																					<b>mfcrr</b>		
<b>mfcrr</b>	31 (0x1F)					rD			///					0	0	0	0	0	1	0	0	1	1	/	X	<b>mfcrr</b>								
<b>mfddcr</b>	31 (0x1F)					rD			DCRN5–9				DCRN0–4				0	1	0	1	0	0	0	0	1	1	/	XF	<b>mfddcr</b>					
<b>mfmsr<sup>(7)</sup></b>	31 (0x1F)					rD			///					0	0	0	1	0	1	0	0	1	1	/	X	<b>mfmsr</b>								
<b>mfpmr</b>	31 (0x1F)					rD			PMRN5–9				PMRN0–4				0	1	0	1	0	0	1	1	1	0	0	XF	<b>mfpmr</b>					
<b>mfregname</b>	mfregname rD					equivalent to					mfspr rD,SPRn																					<b>mfregname</b>		
<b>mfspir<sup>(11)</sup></b>	31 (0x1F)					rD			SPR[5–9]				SPR[0–4]				0	1	0	1	0	1	0	0	1	1	/	XF	<b>mfspir</b>					
<b>mr</b>	mr rA,rS					equivalent to					or rA,rS,rS																					<b>mr</b>		
<b>msync</b>	31 (0x1F)					///					1	0	0	1	0	1	0	1	1	0	/	X	<b>msync</b>											
<b>mtcrr</b>	mtcrr rS					equivalent to					mtcrr 0xFF,rS																					<b>mtcrr</b>		
<b>mtcrf</b>	31 (0x1F)					rS		/		CRM				/		0	0	1	0	0	1	0	0	0	0	0	/	XF	<b>mtcrf</b>					
<b>mtddcr</b>	31 (0x1F)					rS			DCRN5–9				DCRN0–4				0	1	1	1	0	0	0	0	1	1	/	XF	<b>mtddcr</b>					
<b>mtmsr<sup>(7)</sup></b>	31 (0x1F)					rS			///					0	0	1	0	0	1	0	0	1	0	/	X	<b>mtmsr</b>								
<b>mtpmr</b>	31 (0x1F)					rS			PMRN5–9				PMRN0–4				0	1	1	1	0	0	1	1	1	0	0	XF	<b>mtpmr</b>					
<b>mtregname</b>	mtregname rS					equivalent to					mtspir SPRn rS																					<b>mtregname</b>		
<b>mtspir<sup>(11)</sup></b>	31 (0x1F)					rS			SPR[5–9]				SPR[0–4]				0	1	1	1	0	1	0	0	1	1	/	XF	<b>mtspir</b>					
<b>mulhw</b>	31 (0x1F)					rD			rA				rB				/	0	0	1	0	0	1	0	1	1	0	X	<b>mulhw</b>					
<b>mulhw.</b>	31 (0x1F)					rD			rA				rB				/	0	0	1	0	0	1	0	1	1	1	X	<b>mulhw.</b>					
<b>mulhwu</b>	31 (0x1F)					rD			rA				rB				/	0	0	0	0	0	1	0	1	1	0	X	<b>mulhwu</b>					
<b>mulhwu.</b>	31 (0x1F)					rD			rA				rB				/	0	0	0	0	0	1	0	1	1	1	X	<b>mulhwu.</b>					



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>mulli</b>	07						rD			rA			SIMM															D	<b>mulli</b>					
<b>mullw</b>	31 (0x1F)						rD			rA			rB			0	0	1	1	1	0	1	0	1	0	1	1	0	X	<b>mullw</b>				
<b>mullw.</b>	31 (0x1F)						rD			rA			rB			0	0	1	1	1	0	1	0	1	0	1	1	1	X	<b>mullw.</b>				
<b>mullwo</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	1	0	1	0	1	0	1	1	0	X	<b>mullwo</b>				
<b>mullwo.</b>	31 (0x1F)						rD			rA			rB			1	0	1	1	1	0	1	0	1	0	1	1	1	X	<b>mullwo.</b>				
<b>nand</b>	31 (0x1F)						rS			rA			rB			0	1	1	1	0	1	1	1	0	0	0	X	<b>nand</b>						
<b>nand.</b>	31 (0x1F)						rS			rA			rB			0	1	1	1	0	1	1	1	0	0	1	X	<b>nand.</b>						
<b>neg</b>	31 (0x1F)						rD			rA			///			0	0	0	1	1	0	1	0	0	0	0	X	<b>neg</b>						
<b>neg.</b>	31 (0x1F)						rD			rA			///			0	0	0	1	1	0	1	0	0	0	1	X	<b>neg.</b>						
<b>nego</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	1	0	0	0	0	X	<b>nego</b>						
<b>nego.</b>	31 (0x1F)						rD			rA			///			1	0	0	1	1	0	1	0	0	0	1	X	<b>nego.</b>						
<b>nop</b>	nop						equivalent to						ori 0,0,0												<b>nop</b>									
<b>nor</b>	31 (0x1F)						rS			rA			rB			0	0	0	1	1	1	1	1	0	0	0	X	<b>nor</b>						
<b>nor.</b>	31 (0x1F)						rS			rA			rB			0	0	0	1	1	1	1	1	0	0	1	X	<b>nor.</b>						
<b>not</b>	not rA,rS						equivalent to						nor rA,rS,rS												<b>not</b>									
<b>or</b>	31 (0x1F)						rS			rA			rB			0	1	1	0	1	1	1	1	0	0	0	X	<b>or</b>						
<b>or.</b>	31 (0x1F)						rS			rA			rB			0	1	1	0	1	1	1	1	0	0	1	X	<b>or.</b>						
<b>orc</b>	31 (0x1F)						rS			rA			rB			0	1	1	0	0	1	1	1	0	0	0	X	<b>orc</b>						
<b>orc.</b>	31 (0x1F)						rS			rA			rB			0	1	1	0	0	1	1	1	0	0	1	X	<b>orc.</b>						
<b>ori</b>	24 (0x18)						rS			rA			UIMM															D	<b>ori</b>					
<b>oris</b>	25 (0x19)						rS			rA			UIMM															D	<b>oris</b>					
<b>rfci</b>	19 (0x13)						///															0	0	0	0	1	1	0	0	1	1	/	XL	<b>rfci</b>
<b>rfdi</b> <sup>(7)</sup>	0	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	1	1	0	X	<b>rfdi</b>	
<b>rfi</b> <sup>(7)</sup>	19 (0x13)						///															0	0	0	0	1	1	0	0	1	0	/	XL	<b>rfi</b>
<b>rfmci</b> <sup>(7)</sup>	19 (0x13)						///															0	0	0	0	1	0	0	1	1	0	/	XL	<b>rfmci</b>



**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic					
<b>rlwimi</b>	20 (0x14)				rS				rA				SH				MB				ME				0	M	<b>rlwimi</b>												
<b>rlwimi.</b>	20 (0x14)				rS				rA				SH				MB				ME				1	M	<b>rlwimi.</b>												
<b>rlwinm</b>	21 (0x15)				rS				rA				SH				MB				ME				0	M	<b>rlwinm</b>												
<b>rlwinm.</b>	21 (0x15)				rS				rA				SH				MB				ME				1	M	<b>rlwinm.</b>												
<b>rlwnm</b>	23 (0x17)				rS				rA				rB				MB				ME				0	M	<b>rlwnm</b>												
<b>rlwnm.</b>	23 (0x17)				rS				rA				rB				MB				ME				1	M	<b>rlwnm.</b>												
<b>rotlw</b>	rotlw rA,rS,rB				equivalent to				rlwnm rA,rS,rB,0,31																						<b>rotlw</b>								
<b>rotlwi</b>	rotlwi rA,rS,n				equivalent to				rlwinm rA,rS,n,0,31																										<b>rotlwi</b>				
<b>rotrwi</b>	rotrwi rA,rS,n				equivalent to				rlwinm rA,rS,32 – n,0,31																														<b>rotrwi</b>
<b>sc</b>	17 (0x11)								///																1	/	SC			<b>sc</b>									
<b>slw</b>	31 (0x1F)				rS				rA				rB				0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	X	<b>slw</b>					
<b>slw.</b>	31 (0x1F)				rS				rA				rB				0	0	0	0	0	0	1	1	0	0	0	0	1	X	<b>slw.</b>								
<b>slwi</b>	slwi rA,rS,n (n < 32)				equivalent to				rlwinm rA,rS,n,0,31 – n																										<b>slwi</b>				
<b>sraw</b>	31 (0x1F)				rS				rA				rB				1	1	0	0	0	0	1	1	0	0	0	0	0	0	X	<b>sraw</b>							
<b>sraw.</b>	31 (0x1F)				rS				rA				rB				1	1	0	0	0	1	1	0	0	0	0	1	X	<b>sraw.</b>									
<b>srawi</b>	31 (0x1F)				rS				rA				SH				1	1	0	0	1	1	1	0	0	0	0	0	X	<b>srawi</b>									
<b>srawi.</b>	31 (0x1F)				rS				rA				SH				1	1	0	0	1	1	1	0	0	0	1	X	<b>srawi.</b>										
<b>srw</b>	31 (0x1F)				rS				rA				rB				1	0	0	0	0	1	1	0	0	0	0	X	<b>srw</b>										
<b>srw.</b>	31 (0x1F)				rS				rA				rB				1	0	0	0	0	1	1	0	0	0	1	X	<b>srw.</b>										
<b>srwi</b>	srwi rA,rS,n (n < 32)				equivalent to				rlwinm rA,rS,32 – n,n,31																										<b>srwi</b>				
<b>stb</b>	38(0x26)				rS				rA				D																				D	<b>stb</b>					
<b>stbu</b>	39(0x27)				rS				rA				D																								D	<b>stbu</b>	
<b>stbux</b>	31 (0x1F)				rS				rA				rB				0	0	1	1	1	1	0	1	1	1	0	X	<b>stbux</b>										
<b>stbx</b>	31 (0x1F)				rS				rA				rB				0	0	1	1	0	1	0	1	1	1	0	X	<b>stbx</b>										
<b>sth</b>	44(0x2C)				rS				rA				D																				D	<b>sth</b>					



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>sthbrx</b>	31 (0x1F)			rS			rA			rB			1	1	1	0	0	1	0	1	1	0	/	X	<b>sthbrx</b>									
<b>sthu</b>	45(0x2D)			rS			rA			D											D	<b>sthu</b>												
<b>sthux</b>	31 (0x1F)			rS			rA			rB			0	1	1	0	1	1	0	1	1	1	/	X	<b>sthux</b>									
<b>sthx</b>	31 (0x1F)			rS			rA			rB			0	1	1	0	0	1	0	1	1	1	/	X	<b>sthx</b>									
<b>stmw</b>	47(0x2F)			rS			rA			D											D	<b>stmw</b>												
<b>stw</b>	36(0x24)			rS			rA			D											D	<b>stw</b>												
<b>stwbrx</b>	31 (0x1F)			rS			rA			rB			1	0	1	0	0	1	0	1	1	0	/	X	<b>stwbrx</b>									
<b>stwcx.</b>	31 (0x1F)			rS			rA			rB			0	0	1	0	0	1	0	1	1	0	1	X	<b>stwcx.</b>									
<b>stwu</b>	37(0x25)			rS			rA			D											D	<b>stwu</b>												
<b>stwux</b>	31 (0x1F)			rS			rA			rB			0	0	1	0	1	1	0	1	1	1	/	D	<b>stwux</b>									
<b>stwx</b>	31 (0x1F)			rS			rA			rB			0	0	1	0	0	1	0	1	1	1	/	D	<b>stwx</b>									
<b>sub</b>	sub rD,rA,rB			equivalent to			subf rD,rB,rA															<b>sub</b>												
<b>subc</b>	subc rD,rA,rB			equivalent to			subfc rD,rB,rA															<b>subc</b>												
<b>subf</b>	31 (0x1F)			rD			rA			rB			0	0	0	0	1	0	1	0	0	0	0	X	<b>subf</b>									
<b>subf.</b>	31 (0x1F)			rD			rA			rB			0	0	0	0	1	0	1	0	0	0	1	X	<b>subf.</b>									
<b>subfc</b>	31 (0x1F)			rD			rA			rB			0	0	0	0	0	0	1	0	0	0	0	X	<b>subfc</b>									
<b>subfc.</b>	31 (0x1F)			rD			rA			rB			0	0	0	0	0	0	1	0	0	0	1	X	<b>subfc.</b>									
<b>subfco</b>	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	1	0	0	0	0	X	<b>subfco</b>									
<b>subfco.</b>	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	1	0	0	0	1	X	<b>subfco.</b>									
<b>subfe</b>	31 (0x1F)			rD			rA			rB			0	0	1	0	0	0	1	0	0	0	0	X	<b>subfe</b>									
<b>subfe.</b>	31 (0x1F)			rD			rA			rB			0	0	1	0	0	0	1	0	0	0	1	X	<b>subfe.</b>									
<b>subfeo</b>	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	1	0	0	0	0	X	<b>subfeo</b>									
<b>subfeo.</b>	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	1	0	0	0	1	X	<b>subfeo.</b>									
<b>subfic</b>	08			rD			rA			SIMM											D	<b>subfic</b>												
<b>subfme</b>	31 (0x1F)			rD			rA			///			0	0	1	1	1	0	1	0	0	0	0	X	<b>subfme</b>									





Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>subfme.</b>	31 (0x1F)			rD			rA			///			0	0	1	1	1	0	1	0	0	0	0	1	X	<b>subfme.</b>								
<b>subfmeo</b>	31 (0x1F)			rD			rA			///			1	0	1	1	1	0	1	0	0	0	0	0	X	<b>subfmeo</b>								
<b>subfmeo.</b>	31 (0x1F)			rD			rA			///			1	0	1	1	1	0	1	0	0	0	1	X	<b>subfmeo.</b>									
<b>subfo</b>	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	1	0	0	0	0	X	<b>subfo</b>									
<b>subfo.</b>	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	1	0	0	0	1	X	<b>subfo.</b>									
<b>subfze</b>	31 (0x1F)			rD			rA			///			0	0	1	1	0	0	1	0	0	0	0	X	<b>subfze</b>									
<b>subfze.</b>	31 (0x1F)			rD			rA			///			0	0	1	1	0	0	1	0	0	0	1	X	<b>subfze.</b>									
<b>subfzeo</b>	31 (0x1F)			rD			rA			///			1	0	1	1	0	0	1	0	0	0	0	X	<b>subfzeo</b>									
<b>subfzeo.</b>	31 (0x1F)			rD			rA			///			1	0	1	1	0	0	1	0	0	0	1	X	<b>subfzeo.</b>									
<b>subi</b>	subi rD,rA,value			equivalent to			addi rD,rA,-value																				<b>subi</b>							
<b>subic</b>	subic rD,rA,value			equivalent to			addic rD,rA,-value																				<b>subic</b>							
<b>subic.</b>	subic. rD,rA,value			equivalent to			addic. rD,rA,-value																				<b>subic.</b>							
<b>subis</b>	subis rD,rA,value			equivalent to			addis rD,rA,-value																				<b>subis</b>							
<b>tlbie</b> <sup>(6)(7)</sup>	31 (0x1F)			///			///			rB			0	1	0	0	1	1	0	0	1	0	0	X	<b>tlbie</b>									
<b>tlbivax</b>	31 (0x1F)			///			rA			rB			1	1	0	0	0	1	0	0	1	0	/	X	<b>tlbivax</b>									
<b>tlbre</b>	31 (0x1F)			/// <sup>(12)</sup>									1	1	1	0	1	1	0	0	1	0	/	X	<b>tlbre</b>									
<b>tlbsx</b>	31 (0x1F)			/// <sup>(12)</sup>			rA			rB			1	1	1	0	0	1	0	0	1	0	/12	X	<b>tlbsx</b>									
<b>tlbsync</b> <sup>(6)(7)</sup>	31 (0x1F)			///			///			///			1	0	0	0	1	1	0	1	1	0	/	X	<b>tlbsync</b>									
<b>tlbwe</b>	31 (0x1F)			/// <sup>(12)</sup>									1	1	1	1	0	1	0	0	1	0	/	X	<b>tlbwe</b>									
<b>tw</b>	31 (0x1F)			TO			rA			rB			0	0	0	0	0	0	0	1	0	0	/	X	<b>tw</b>									
<b>tweq</b>	tweq rA,SIMM			equivalent to			tw 4,rA,SIMM																				<b>tweq</b>							
<b>tweqi</b>	tweqi rA,SIMM			equivalent to			twi 4,rA,SIMM																				<b>tweqi</b>							
<b>twge</b>	twge rA,SIMM			equivalent to			tw 12,rA,SIMM																				<b>twge</b>							
<b>twgei</b>	twgei rA,SIMM			equivalent to			twi 12,rA,SIMM																				<b>twgei</b>							



Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>twgt</b>	twgt rA,SIMM		equivalent to		tw 8,rA,SIMM																						<b>twgt</b>							
<b>twgti</b>	twgti rA,SIMM		equivalent to		twi 8,rA,SIMM																						<b>twgti</b>							
<b>twi</b>	03			TO			rA			SIMM												D			<b>twi</b>									
<b>twle</b>	twle rA,SIMM		equivalent to		tw 20,rA,SIMM																						<b>twle</b>							
<b>twlei</b>	twlei rA,SIMM		equivalent to		twi 20,rA,SIMM																						<b>twlei</b>							
<b>twlge</b>	twlge rA,SIMM		equivalent to		tw 12,rA,SIMM																						<b>twlge</b>							
<b>twlgei</b>	twlgei rA,SIMM		equivalent to		twi 12,rA,SIMM																						<b>twlgei</b>							
<b>twlgt</b>	twlgt rA,SIMM		equivalent to		tw 1,rA,SIMM																						<b>twlgt</b>							
<b>twlgti</b>	twlgti rA,SIMM		equivalent to		twi 1,rA,SIMM																						<b>twlgti</b>							
<b>twlle</b>	twlle rA,SIMM		equivalent to		tw 6,rA,SIMM																						<b>twlle</b>							
<b>twllei</b>	twllei rA,SIMM		equivalent to		twi 6,rA,SIMM																						<b>twllei</b>							
<b>twllt</b>	twllt rA,SIMM		equivalent to		tw 2,rA,SIMM																						<b>twllt</b>							
<b>twllti</b>	twllti rA,SIMM		equivalent to		twi 2,rA,SIMM																						<b>twllti</b>							
<b>twlng</b>	twlng rA,SIMM		equivalent to		tw 6,rA,SIMM																						<b>twlng</b>							
<b>twlngi</b>	twlngi rA,SIMM		equivalent to		twi 6,rA,SIMM																						<b>twlngi</b>							
<b>twlnl</b>	twlnl rA,SIMM		equivalent to		tw 5,rA,SIMM																						<b>twlnl</b>							
<b>twlnli</b>	twlnli rA,SIMM		equivalent to		twi 5,rA,SIMM																						<b>twlnli</b>							
<b>twlt</b>	twlt rA,SIMM		equivalent to		tw 16,rA,SIMM																						<b>twlt</b>							
<b>twlti</b>	twlti rA,SIMM		equivalent to		twi 16,rA,SIMM																						<b>twlti</b>							
<b>twne</b>	twne rA,SIMM		equivalent to		tw 24,rA,SIMM																						<b>twne</b>							
<b>twnei</b>	twnei rA,SIMM		equivalent to		twi 24,rA,SIMM																						<b>twnei</b>							
<b>twng</b>	twng rA,SIMM		equivalent to		tw 20,rA,SIMM																						<b>twng</b>							
<b>twngi</b>	twngi rA,SIMM		equivalent to		twi 20,rA,SIMM																						<b>twngi</b>							
<b>twnl</b>	twnl rA,SIMM		equivalent to		tw 12,rA,SIMM																						<b>twnl</b>							
<b>twnli</b>	twnli rA,SIMM		equivalent to		twi 12,rA,SIMM																						<b>twnli</b>							

**Table 269. Instructions sorted by mnemonic (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic										
<b>wait</b>	31 (0x1F)						///															0	0	0	0	1	1	1	1	1	1	0	/			wait								
<b>wrtee</b>	31 (0x1F)						rS			///															0	0	1	0	0	0	0	0	0	1	1	/	X		wrtee					
<b>wrteei</b>	31 (0x1F)						///						E	///															0	0	1	0	1	0	0	0	1	1	/	X		wrteei		
<b>xor</b>	31 (0x1F)						rS			rA			rB						0	1	0	0	1	1	1	1	1	0	0	0	X		<b>xor</b>											
<b>xor.</b>	31 (0x1F)						rS			rA			rB						0	1	0	0	1	1	1	1	0	0	1	X		<b>xor.</b>												
<b>xori</b>	26 (0x1A)						rS			rA			UIMM																													D	<b>xori</b>	
<b>xoris</b>	27 (0x1B)						rS			rA			UIMM																														D	<b>xoris</b>

1. Simplified mnemonics for branch instructions that do not test a CR bit should not specify one; a programming error may occur.
2. The value in the BI operand selects CRn[2], the EQ bit.
3. The value in the BI operand selects CRn[0], the LT bit.
4. The value in the BI operand selects CRn[1], the GT bit.
5. The value in the BI operand selects CRn[3], the SO bit.
6. Optional to the PowerPC classic architecture.
7. Supervisor-level instruction.
8. d = UIMM \* 8.
9. d = UIMM \* 2.
10. d = UIMM \* 4.
11. Access level is determined by whether the SPR is defined as a user- or supervisor-level SPR.
12. This field is defined as allocated by the Book E architecture, for possible use in an implementation.

## A.2 Instructions sorted by primary opcodes (decimal and hexadecimal)

[Table 270](#) lists instructions by their primary (0–5) opcodes in decimal and hexadecimal format.





Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic									
efsadd	04				rD				rA				rB				0	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	EFX	efsadd								
efscfd	04				rD				0	0	0	0	0	rB				0	1	0	1	1	0	0	1	1	1	1	1	1	1	1	1	1	EFX	efscfd							
efscfsf	04				rD				///				rB				0	1	0	1	1	0	1	0	0	1	0	0	1	1	1	1	1	1	EFX	efscfsf							
efscfsi	04				rD				///				rB				0	1	0	1	1	0	1	0	0	0	0	0	0	0	0	1	1	1	1	EFX	efscfsi						
efscfuf	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	0	1	0	0	1	0	0	0	0	EFX	efscfuf						
efscfui	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	EFX	efscfui					
efscmpeq	04				crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	1	1	0	1	1	1	0	0	0	0	EFX	efscmpeq						
efscmpgt	04				crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	0	0	1	1	0	0	0	0	0	0	0	EFX	efscmpgt					
efscmplt	04				crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	0	0	1	1	0	1	0	1	0	1	0	1	EFX	efscmplt				
efscstf	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	1	0	1	1	1	1	1	1	1	1	EFX	efscstf					
efscstsi	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	1	0	1	0	0	1	0	1	0	1	0	1	EFX	efscstsi			
efscstsiz	04				rD				///				rB				0	1	0	1	1	0	1	1	0	1	1	0	1	0	1	0	1	0	0	0	0	EFX	efscstsiz				
efscstuf	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	1	0	1	1	1	0	0	0	0	0	0	0	EFX	efscstuf			
efscstui	04				rD				///				rB				0	1	0	1	1	0	1	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	EFX	efscstui		
efscstuiz	04				rD				///				rB				0	1	0	1	1	0	1	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	EFX	efscstuiz		
efsddiv	04				rD				rA				rB				0	1	0	1	1	0	0	1	0	0	1	0	0	0	1	0	0	1	0	0	0	0	EFX	efsddiv			
efsmul	04				rD				rA				rB				0	1	0	1	1	0	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	EFX	efsmul		
efsnabs	04				rD				rA				///				0	1	0	1	1	0	0	0	0	0	1	0	1	0	1	0	1	0	1	0	1	0	1	EFX	efsnabs		
efsneg	04				rD				rA				///				0	1	0	1	1	0	0	0	0	0	1	1	0	0	1	1	0	0	0	0	0	0	0	EFX	efsneg		
efssub	04				rD				rA				rB				0	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0	1	EFX	efssub	
efststeq	04				crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	EFX	efststeq	
efststgt	04				crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	EFX	efststgt
efststlt	04				crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	0	1	1	1	1	1	1	0	1	0	1	0	1	0	1	0	1	EFX	efststlt
mulli	07				rD				rA				SIMM																D	mulli													
subfic	08				rD				rA				SIMM																D	subfic													
cmpli	10 (0x0A)				crfD		/	L	rA				UIMM																D	cmpli													
cmpi	11 (0x0B)				crfD		/	L	rA				SIMM																D	cmpi													
addic	12 (0x0C)				rD				rA				SIMM																D	addic													
addic.	13 (0x0D)				rD				rA				SIMM																D	addic.													



**Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic									
addi	14 (0x0E)				rD				rA				SIMM												D	addi																	
addis	15 (0x0F)				rD				rA				SIMM												D	addis																	
bc	16 (0x10)				BO				BI				BD												0	0	B	bc															
bca	16 (0x10)				BO				BI				BD												1	0	B	bca															
bcl	16 (0x10)				BO				BI				BD												0	1	B	bcl															
bcla	16 (0x10)				BO				BI				BD												1	1	B	bcla															
sc	17 (0x11)				///												1	/	SC	sc																							
b	18 (0x12)				LI												0	0	I	b																							
ba	18 (0x12)				LI												1	0	I	ba																							
bl	18 (0x12)				LI												0	1	I	bl																							
bla	18 (0x12)				LI												1	1	I	bla																							
rfci	19 (0x13)				///												0	0	0	0	1	1	0	0	1	1	/	XL	rfci														
rfmci <sup>(1)</sup>	19 (0x13)				///												0	0	0	0	1	0	0	1	1	0	/	XL	rfmci														
mcrf	19 (0x13)				crfD	//	crfS				///												0	0	0	0	0	0	0	0	0	0	0	0	0	/	XL	mcrf					
bclr	19 (0x13)				BO				BI				///												0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	XL	bclr		
bclrl	19 (0x13)				BO				BI				///												0	0	0	0	0	1	0	0	0	0	0	1	0	0	0	0	XL	bclrl	
crnor	19 (0x13)				crbD				crbA				crbB				0	0	0	0	1	0	0	0	0	0	0	0	1	/	XL	crnor											
rfi <sup>(1)</sup>	19 (0x13)				///												0	0	0	0	1	1	0	0	1	0	/	XL	rfi														
crandc	19 (0x13)				crbD				crbA				crbB				0	0	1	0	0	0	0	0	0	0	0	0	1	/	XL	crandc											
isync	19 (0x13)				///												0	0	1	0	0	1	0	1	1	0	/	XL	isync														
crxor	19 (0x13)				crbD				crbA				crbB				0	0	1	1	0	0	0	0	0	0	0	0	1	/	XL	crxor											
crand	19 (0x13)				crbD				crbA				crbB				0	1	0	0	0	0	0	0	0	0	0	1	/	XL	crand												
crnand	19 (0x13)				crbD				crbA				crbB				0	0	1	1	1	0	0	0	0	0	1	/	XL	crnand													
creqv	19 (0x13)				crbD				crbA				crbB				0	1	0	0	1	0	0	0	0	0	1	/	XL	creqv													
crorc	19 (0x13)				crbD				crbA				crbB				0	1	1	0	1	0	0	0	0	0	1	/	XL	crorc													
cror	19 (0x13)				crbD				crbA				crbB				0	1	1	1	0	0	0	0	0	0	1	/	XL	cror													
bcctr	19 (0x13)				BO				BI				///												1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	XL	bcctr	
bcctrl	19 (0x13)				BO				BI				///												1	0	0	0	0	1	0	0	0	0	0	1	0	0	0	0	1	XL	bcctrl



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
rlwimi	20 (0x14)				rS				rA				SH				MB				ME				0	M	rlwimi							
rlwimi.	20 (0x14)				rS				rA				SH				MB				ME				1	M	rlwimi.							
rlwinm	21 (0x15)				rS				rA				SH				MB				ME				0	M	rlwinm							
rlwinm.	21 (0x15)				rS				rA				SH				MB				ME				1	M	rlwinm.							
rlwnm	23 (0x17)				rS				rA				rB				MB				ME				0	M	rlwnm							
rlwnm.	23 (0x17)				rS				rA				rB				MB				ME				1	M	rlwnm.							
ori	24 (0x18)				rS				rA				UIMM												D	ori								
oris	25 (0x19)				rS				rA				UIMM												D	oris								
xori	26 (0x1A)				rS				rA				UIMM												D	xori								
xoris	27 (0x1B)				rS				rA				UIMM												D	xoris								
andi.	28 (0x1C)				rS				rA				UIMM												D	andi.								
andis.	29 (0x1D)				rS				rA				UIMM												D	andis.								
dcblc	31 (0x1F)				CT				rA				rB				0	1	1	0	0	0	0	0	0	1	1	0	0	X	dcblc			
dcbtls	31 (0x1F)				CT				rA				rB				0	0	1	0	1	0	0	1	1	0	0	X	dcbtls					
dcbstls	31 (0x1F)				CT				rA				rB				0	0	1	0	0	0	0	1	1	0	0	X	dcbstls					
evabs	31 (0x1F)				rD				rA				///				0	1	0	0	0	0	0	1	0	0	0	EVX	evabs					
evaddiw	31 (0x1F)				rD				UIMM				rB				0	1	0	0	0	0	0	0	0	1	0	EVX	evaddiw					
evaddsmiaaw	31 (0x1F)				rD				rA				///				1	0	0	1	1	0	0	1	0	0	1	EVX	evaddsmiaaw					
evaddssiaaw	31 (0x1F)				rD				rA				///				1	0	0	1	1	0	0	0	0	0	1	EVX	evaddssiaaw					
evaddumiaaw	31 (0x1F)				rD				rA				///				1	0	0	1	1	0	0	1	0	0	0	EVX	evaddumiaaw					
evaddusiaaw	31 (0x1F)				rD				rA				///				1	0	0	1	1	0	0	0	0	0	0	EVX	evaddusiaaw					
evaddw	31 (0x1F)				rD				rA				rB				0	1	0	0	0	0	0	0	0	0	0	0	EVX	evaddw				
evand	31 (0x1F)				rD				rA				rB				0	1	0	0	0	0	1	0	0	0	1	EVX	evand					
evandc	31 (0x1F)				rD				rA				rB				0	1	0	0	0	0	1	0	0	1	0	EVX	evandc					
evcmpeq	31 (0x1F)				crfD				/	/	rA				rB				0	1	0	0	0	1	1	0	1	0	0	EVX	evcmpeq			
evcmpgts	31 (0x1F)				crfD				/	/	rA				rB				0	1	0	0	0	1	1	0	0	0	1	EVX	evcmpgts			
evcmpgtu	31 (0x1F)				crfD				/	/	rA				rB				0	1	0	0	0	1	1	0	0	0	0	EVX	evcmpgtu			
evcmplt	31 (0x1F)				crfD				/	/	rA				rB				0	1	0	0	0	1	1	0	0	1	1	EVX	evcmplt			
evcmpltu	31 (0x1F)				crfD				/	/	rA				rB				0	1	0	0	0	1	1	0	0	1	0	EVX	evcmpltu			





Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evcntlsw	31 (0x1F)			rD					rA					///					0	1	0	0	0	0	0	0	1	1	1	0	EVX	evcntlsw		
evcntlzw	31 (0x1F)			rD					rA					///					0	1	0	0	0	0	0	0	1	1	0	1	EVX	evcntlzw		
evdivws	31 (0x1F)			rD					rA					rB					1	0	0	1	1	0	0	0	0	1	1	0	EVX	evdivws		
evdivwu	31 (0x1F)			rD					rA					rB					1	0	0	1	1	0	0	0	0	1	1	1	EVX	evdivwu		
eveqv	31 (0x1F)			rD					rA					rB					0	1	0	0	0	0	1	1	0	0	1	EVX	eveqv			
evextsb	31 (0x1F)			rD					rA					///					0	1	0	0	0	0	0	1	0	1	0	EVX	evextsb			
evextsh	31 (0x1F)			rD					rA					///					0	1	0	0	0	0	0	1	0	1	1	EVX	evextsh			
evfsabs	31 (0x1F)			rD					rA					///					0	1	0	1	0	0	0	0	1	0	0	EVX	evfsabs			
evfsadd	31 (0x1F)			rD					rA					rB					0	1	0	1	0	0	0	0	0	0	0	0	EVX	evfsadd		
evfscfsf	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	0	1	1	EVX	evfscfsf			
evfscfsi	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	0	0	1	EVX	evfscfsi			
evfscfuf	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	0	1	0	EVX	evfscfuf			
evfscfui	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	0	0	0	EVX	evfscfui			
evfscmpeq	31 (0x1F)			crfD		/	/	rA					rB					0	1	0	1	0	0	0	1	1	1	0	EVX	evfscmpeq				
evfscmpgt	31 (0x1F)			crfD		/	/	rA					rB					0	1	0	1	0	0	0	1	1	0	0	EVX	evfscmpgt				
evfscmplt	31 (0x1F)			crfD		/	/	rA					rB					0	1	0	1	0	0	0	1	1	0	1	EVX	evfscmplt				
evfsctsf	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	1	1	1	EVX	evfsctsf			
evfsctsi	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	1	0	1	EVX	evfsctsi			
evfsctsiz	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	1	0	1	0	EVX	evfsctsiz			
evfsctuf	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	1	1	0	EVX	evfsctuf			
evfsctui	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	0	1	0	0	EVX	evfsctui			
evfsctuiz	31 (0x1F)			rD					///					rB					0	1	0	1	0	0	1	1	0	0	0	EVX	evfsctuiz			
evfsdiv	31 (0x1F)			rD					rA					rB					0	1	0	1	0	0	0	1	0	0	1	EVX	evfsdiv			
evfsmul	31 (0x1F)			rD					rA					rB					0	1	0	1	0	0	0	1	0	0	0	EVX	evfsmul			
evfsnabs	31 (0x1F)			rD					rA					///					0	1	0	1	0	0	0	0	1	0	1	EVX	evfsnabs			
evfsneg	31 (0x1F)			rD					rA					///					0	1	0	1	0	0	0	0	1	1	0	EVX	evfsneg			
evfssub	31 (0x1F)			rD					rA					rB					0	1	0	1	0	0	0	0	0	0	1	EVX	evfssub			
evfststeq	31 (0x1F)			crfD		/	/	rA					rB					0	1	0	1	0	0	1	1	1	1	0	EVX	evfststeq				
evfststgt	31 (0x1F)			crfD		/	/	rA					rB					0	1	0	1	0	0	1	1	1	0	0	EVX	evfststgt				



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evfststlt	31 (0x1F)			crfD		/	/	rA				rB				0	1	0	1	0	0	1	1	1	0	1	EVX	evfststlt						
evlidd	31 (0x1F)			rD				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	EVX	evlidd		
evlddx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evlddx	
evldh	31 (0x1F)			rD				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	0	0	0	0	1	0	1	EVX	evldh					
evldhx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	0	0	1	0	0	EVX	evldhx					
evldw	31 (0x1F)			rD				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	0	0	0	0	0	1	1	EVX	evldw					
evldwx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	0	0	0	1	0	EVX	evldwx					
evlhhesplat	31 (0x1F)			rD				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	0	0	1	0	0	1	EVX	evlhhesplat						
evlhhesplatx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	1	0	0	0	EVX	evlhhesplatx						
evlhhosplat	31 (0x1F)			rD				rA				UIMM <sup>(3)</sup>				0	1	1	0	0	0	0	1	1	1	1	EVX	evlhhosplat						
evlhhosplatx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	1	1	1	0	EVX	evlhhosplatx						
evlhhusplat	31 (0x1F)			rD				rA				UIMM <sup>(3)</sup>				0	1	1	0	0	0	0	1	1	0	1	EVX	evlhhusplat						
evlhhusplatx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	0	1	1	0	0	EVX	evlhhusplatx						
evlwhe	31 (0x1F)			rD				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	0	1	0	0	0	1	EVX	evlwhe						
evlwhex	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	1	0	0	0	0	EVX	evlwhex						
evlwhos	31 (0x1F)			rD				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	0	1	0	1	1	1	EVX	evlwhos						
evlwhosx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	1	0	1	1	0	EVX	evlwhosx						
evlwhou	31 (0x1F)			rD				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	0	1	0	1	0	1	EVX	evlwhou						
evlwhoux	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	1	0	1	0	0	EVX	evlwhoux						
evlwhsplat	31 (0x1F)			rD				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	0	1	1	1	0	1	EVX	evlwhsplat						
evlwhsplatx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	1	1	1	0	0	EVX	evlwhsplatx						
evlwwsplat	31 (0x1F)			rD				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	0	1	1	0	0	1	EVX	evlwwsplat						
evlwwsplatx	31 (0x1F)			rD				rA				rB				0	1	1	0	0	0	1	1	0	0	0	EVX	evlwwsplatx						
evmergehi	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	1	1	0	0	EVX	evmergehi						
evmergehilo	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	1	1	1	0	EVX	evmergehilo						
evmergeelo	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	1	1	0	1	EVX	evmergeelo						
evmergelohi	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	1	1	1	1	EVX	evmergelohi						



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhegsmfaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	0	1	0	1	0	1	1	EVX	evmhegsmfaa					
evmhegsmfan	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	1	0	1	1	EVX	evmhegsmfan					
evmhegsmiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	0	0	1	0	0	1	EVX	evmhegsmiaa						
evmhegsmian	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	0	0	1	EVX	evmhegsmian						
evmhegumiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	0	1	0	0	0	0	EVX	evmhegumiaa						
evmhegumian	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	0	0	0	EVX	evmhegumian						
evmhesmf	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	1	0	1	1	1	EVX	evmhesmf								
evmhesmfa	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	0	1	0	1	1	EVX	evmhesmfa							
evmhesmfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	1	0	1	1	1	EVX	evmhesmfaaw								
evmhesmfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	1	0	1	1	1	EVX	evmhesmfanw								
evmhesmi	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	1	0	0	1	EVX	evmhesmi									
evmhesmia	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	0	0	1	EVX	evmhesmia									
evmhesmiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	1	0	0	1	EVX	evmhesmiaaw									
evmhesmianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	1	0	0	1	EVX	evmhesmianw									
evmhessf	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	0	0	1	1	EVX	evmhessf									
evmhessfa	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	0	0	1	1	EVX	evmhessfa									
evmhessfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	1	1	EVX	evmhessfaaw									
evmhessfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	1	1	EVX	evmhessfanw									
evmhessiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	0	1	EVX	evmhessiaaw									
evmhessianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	0	1	EVX	evmhessianw									
evmheumi	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	1	0	0	0	EVX	evmheumi									
evmheumia	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	0	0	0	EVX	evmheumia									
evmheumiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	1	0	0	0	EVX	evmheumiaaw									
evmheumianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	1	0	0	0	EVX	evmheumianw									
evmheusiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	EVX	evmheusiaaw								
evmheusianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	EVX	evmheusianw								
evmhogsmfaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	1	1	1	1	EVX	evmhogsmfaa								
evmhogsmfan	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	1	1	1	1	EVX	evmhogsmfan								
evmhogsmiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	1	0	1	EVX	evmhogsmiaa									



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhogsmian	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	1	1	0	1	0	1	EVX	evmhogsmian			
evmhogumiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	0	1	0	1	1	0	0	1	1	0	0	1	0	0	EVX	evmhogumiaa		
evmhogumian	31 (0x1F)			rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	1	1	0	0	1	0	0	EVX	evmhogumian		
evmhosmf	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhosmf		
evmhosmfa	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	0	1	1	1	1	1	1	1	1	1	EVX	evmhosmfa		
evmhosmfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhosmfaaw		
evmhosmfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhosmfanw		
evmhosmi	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	0	1	1	0	1	0	1	0	1	0	1	EVX	evmhosmi		
evmhosmia	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	1	0	1	0	1	0	1	0	1	0	1	EVX	evmhosmia	
evmhosmiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	1	1	0	1	0	1	0	1	0	1	EVX	evmhosmiaaw		
evmhosmianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	1	1	0	1	0	1	0	1	0	1	EVX	evmhosmianw		
evmhossf	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhossf	
evmhossfa	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	0	1	1	1	1	1	1	1	1	1	1	1	EVX	evmhossfa	
evmhossfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhossfaaw	
evmhossfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	EVX	evmhossfanw	
evmhossiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	1	0	1	0	1	0	1	0	1	EVX	evmhossiaaw		
evmhossianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	1	0	1	0	1	0	1	0	1	EVX	evmhossianw		
evmhoumi	31 (0x1F)			rD			rA			rB			1	0	0	0	0	0	0	0	1	1	0	0	1	0	0	1	0	0	EVX	evmhoumi		
evmhoumia	31 (0x1F)			rD			rA			rB			1	0	0	0	0	1	0	1	1	0	0	1	0	0	1	0	0	EVX	evmhoumia			
evmhoumiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	1	1	0	0	1	0	0	1	0	0	EVX	evmhoumiaaw		
evmhoumianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	1	1	0	0	1	0	0	1	0	0	EVX	evmhoumianw		
evmhousiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	0	0	0	0	0	1	0	0	1	0	0	1	0	0	EVX	evmhousiaaw		
evmhousianw	31 (0x1F)			rD			rA			rB			1	0	1	1	0	0	0	0	0	1	0	0	1	0	0	1	0	0	EVX	evmhousianw		
evmra	31 (0x1F)			rD			rA			///			1	0	0	1	1	0	0	0	0	1	0	0	1	0	0	1	0	0	EVX	evmra		
evmwhgsmfaa	31 (0x1F)			rD			rA			rB			1	0	1	0	1	1	0	1	1	1	0	1	1	1	1	1	1	1	EVX	evmwhgsmfaa		
evmwhgsmfan	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	1	1	1	1	0	1	1	1	1	1	1	1	EVX	evmwhgsmfan		
evmwhgsmiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	1	1	0	1	1	0	1	1	0	1	1	0	1	0	1	EVX	evmwhgsmiaa	
evmwhgsmian	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	1	1	1	0	1	0	1	EVX	evmwhgsmian		
evmwhgssfaa	31 (0x1F)			rD			rA			rB			1	0	1	0	1	1	0	0	1	1	0	0	1	1	1	1	1	1	EVX	evmwhgssfaa		



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmwhgssfan	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	1	0	1	1	1	1	0	1	0	1	1	1	EVX	evmwhgssfan		
evmwhgumiaa	31 (0x1F)			rD			rA			rB			1	0	1	0	1	1	0	1	1	0	1	1	0	1	1	0	0	EVX	evmwhgumiaa			
evmwhgumian	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	1	1	0	0	EVX	evmwhgumian					
evmwhsmf	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmf				
evmwhsmfa	31 (0x1F)			rD			rA			rB			1	0	0	0	1	1	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfa				
evmwhsmfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfaaw				
evmwhsmfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfanw				
evmwhsmi	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	0	1	1	0	1	1	0	1	EVX	evmwhsmi						
evmwhsmia	31 (0x1F)			rD			rA			rB			1	0	0	0	1	1	0	1	1	0	1	1	0	1	EVX	evmwhsmia						
evmwhsmiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	1	1	0	1	1	0	1	EVX	evmwhsmiaaw						
evmwhsmianw	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	1	0	1	1	0	1	EVX	evmwhsmianw						
evmwhssf	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	0	0	1	1	1	1	1	1	1	EVX	evmwhssf					
evmwhssfa	31 (0x1F)			rD			rA			rB			1	0	0	0	1	1	0	0	1	1	1	1	1	1	1	EVX	evmwhssfa					
evmwhssfaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	0	1	1	1	1	1	1	1	EVX	evmwhssfaaw					
evmwhssfanw	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	0	1	1	1	1	1	1	1	EVX	evmwhssfanw					
evmwhssianw	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	0	1	0	1	0	1	EVX	evmwhssianw							
evmwhssmaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	0	1	0	1	0	1	EVX	evmwhssmaaw							
evmwhumi	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	0	1	1	0	1	1	0	0	EVX	evmwhumi						
evmwhumia	31 (0x1F)			rD			rA			rB			1	0	0	0	1	1	0	1	1	0	1	1	0	0	EVX	evmwhumia						
evmwhusiaaw	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	0	1	0	0	1	0	0	EVX	evmwhusiaaw						
evmwhusianw	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	0	1	0	0	1	0	0	EVX	evmwhusianw						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	0	0	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	0	0	1	1	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	0	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB			1	0	1	1	1	0	0	1	0	1	1	1	1	1	EVX	evmwlsmf						
evmwlsmf	31 (0x1F)			rD			rA			rB</																								



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
evmwlsffaaw	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	0	0	0	0	1	1	EVX	evmwlsffaaw			
evmwlsffanw	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	0	0	0	0	0	0	1	1	EVX	evmwlsffanw	
evmwlsffaaw	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	0	0	0	0	0	0	0	1	EVX	evmwlsffaaw	
evmwlsffanw	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	0	0	0	0	0	0	0	1	EVX	evmwlsffanw	
evmwlsffaaw	31 (0x1F)			rD					rA					rB					1	0	0	0	1	0	0	1	0	0	1	0	0	0	EVX	evmwlsffaaw	
evmwlsffanw	31 (0x1F)			rD					rA					rB					1	0	0	0	1	1	0	1	0	1	0	0	0	0	EVX	evmwlsffanw	
evmwlsffaaw	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	0	1	0	0	1	0	0	0	EVX	evmwlsffaaw	
evmwlsffanw	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	0	1	0	0	1	0	0	0	EVX	evmwlsffanw	
evmwlsffaaw	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	EVX	evmwlsffaaw
evmwlsffanw	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	0	0	0	0	0	0	0	0	0	EVX	evmwlsffanw
evmwsmf	31 (0x1F)			rD					rA					rB					1	0	0	0	1	0	1	1	1	0	1	1	EVX	evmwsmf			
evmwsmfa	31 (0x1F)			rD					rA					rB					1	0	0	0	1	1	1	1	1	1	0	1	1	EVX	evmwsmfa		
evmwsmfaa	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	1	1	0	1	1	EVX	evmwsmfaa				
evmwsmfan	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	1	1	0	1	1	EVX	evmwsmfan				
evmwsmi	31 (0x1F)			rD					rA					rB					1	0	0	0	1	0	1	1	0	0	1	EVX	evmwsmi				
evmwsmia	31 (0x1F)			rD					rA					rB					1	0	0	0	1	1	1	1	0	0	1	EVX	evmwsmia				
evmwsmiaa	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	1	1	0	0	1	EVX	evmwsmiaa				
evmwsmian	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	1	1	0	0	1	EVX	evmwsmian				
evmwssf	31 (0x1F)			rD					rA					rB					1	0	0	0	1	0	1	0	0	1	1	EVX	evmwssf				
evmwssfa	31 (0x1F)			rD					rA					rB					1	0	0	0	1	1	1	0	0	1	1	EVX	evmwssfa				
evmwssfaa	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	1	0	0	1	1	EVX	evmwssfaa				
evmwssfan	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	1	0	0	1	1	EVX	evmwssfan				
evmwumi	31 (0x1F)			rD					rA					rB					1	0	0	0	1	0	1	1	0	0	0	EVX	evmwumi				
evmwumia	31 (0x1F)			rD					rA					rB					1	0	0	0	1	1	1	1	0	0	0	EVX	evmwumia				
evmwumiaa	31 (0x1F)			rD					rA					rB					1	0	1	0	1	0	1	1	0	0	0	EVX	evmwumiaa				
evmwumian	31 (0x1F)			rD					rA					rB					1	0	1	1	1	0	1	1	0	0	0	EVX	evmwumian				
evnand	31 (0x1F)			rD					rA					rB					0	1	0	0	0	0	1	1	1	1	0	EVX	evnand				
evneg	31 (0x1F)			rD					rA					///					0	1	0	0	0	0	0	0	1	0	0	1	EVX	evneg			
evnor	31 (0x1F)			rD					rA					rB					0	1	0	0	0	0	1	1	0	0	0	EVX	evnor				



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evor	31 (0x1F)			rD				rA				rB				0	1	0	0	0	0	0	1	0	1	1	1	1	EVX	evor				
evorc	31 (0x1F)			rD				rA				rB				0	1	0	0	0	0	1	1	0	1	1	EVX	evorc						
evrlw	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	1	0	0	0	EVX	evrlw						
evrlwi	31 (0x1F)			rD				rA				UIMM				0	1	0	0	0	1	0	1	0	1	0	EVX	evrlwi						
evrndw	31 (0x1F)			rD				rA				UIMM				0	1	0	0	0	0	0	1	1	0	0	EVX	evrndw						
evsel	31 (0x1F)			rD				rA				rB				0	1	0	0	1	1	1	1	crfS			EVX	evsel						
evslw	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	0	1	0	0	EVX	evslw						
evslwi	31 (0x1F)			rD				rA				UIMM				0	1	0	0	0	1	0	0	1	1	0	EVX	evslwi						
evsplatfi	31 (0x1F)			rD				SIMM				///				0	1	0	0	0	1	0	1	0	1	1	EVX	evsplatfi						
evsplati	31 (0x1F)			rD				SIMM				///				0	1	0	0	0	1	0	1	0	0	1	EVX	evsplati						
evsrwis	31 (0x1F)			rD				rA				UIMM				0	1	0	0	0	1	0	0	0	1	1	EVX	evsrwis						
evsrwiu	31 (0x1F)			rD				rA				UIMM				0	1	0	0	0	1	0	0	0	1	0	EVX	evsrwiu						
evsrws	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	0	0	0	1	EVX	evsrws						
evsrwu	31 (0x1F)			rD				rA				rB				0	1	0	0	0	1	0	0	0	0	0	EVX	evsrwu						
evstdd	31 (0x1F)			rD				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	1	0	0	0	0	1	EVX	evstdd						
evstddx	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	0	0	0	0	0	EVX	evstddx						
evstdh	31 (0x1F)			rS				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	1	0	0	1	0	1	EVX	evstdh						
evstdhx	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	0	0	1	0	0	EVX	evstdhx						
evstdw	31 (0x1F)			rS				rA				UIMM <sup>(2)</sup>				0	1	1	0	0	1	0	0	0	1	1	EVX	evstdw						
evstdwx	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	0	0	0	1	0	EVX	evstdwx						
evstwhe	31 (0x1F)			rS				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	1	1	0	0	0	1	EVX	evstwhe						
evstwhex	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	1	0	0	0	0	EVX	evstwhex						
evstwho	31 (0x1F)			rS				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	1	1	0	1	0	1	EVX	evstwho						
evstwhox	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	1	0	1	0	0	EVX	evstwhox						
evstwwe	31 (0x1F)			rS				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	1	1	1	0	0	1	EVX	evstwwe						
evstwwex	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	1	1	0	0	0	EVX	evstwwex						
evstwwo	31 (0x1F)			rS				rA				UIMM <sup>(4)</sup>				0	1	1	0	0	1	1	1	1	0	1	EVX	evstwwo						
evstwwox	31 (0x1F)			rS				rA				rB				0	1	1	0	0	1	1	1	1	0	0	EVX	evstwwox						



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evsubfsmiaaw	31 (0x1F)			rD					rA					///					1	0	0	1	1	0	0	1	0	1	1	EVX	evsubfsmiaaw			
evsubfssiaaw	31 (0x1F)			rD					rA					///					1	0	0	1	1	0	0	0	0	0	0	1	1	EVX	evsubfssiaaw	
evsubfumiaaw	31 (0x1F)			rD					rA					///					1	0	0	1	1	0	0	1	0	1	0	EVX	evsubfumiaaw			
evsubfusiaaw	31 (0x1F)			rD					rA					///					1	0	0	1	1	0	0	0	0	0	1	0	EVX	evsubfusiaaw		
evsubfw	31 (0x1F)			rD					rA					rB					0	1	0	0	0	0	0	0	0	1	0	0	EVX	evsubfw		
evsubifw	31 (0x1F)			rD					UIMM					rB					0	1	0	0	0	0	0	0	0	1	1	0	EVX	evsubifw		
evxor	31 (0x1F)			rD					rA					rB					0	1	0	0	0	0	1	0	1	1	0	EVX	evxor			
icblc	31 (0x1F)			CT					rA					rB					0	0	1	1	1	0	0	1	1	0	0	X	icblc			
icbt	31 (0x1F)			CT					rA					rB					0	0	0	0	0	1	0	1	1	0	/	X	icbt			
icbtls	31 (0x1F)			CT					rA					rB					0	1	1	1	1	0	0	1	1	0	0	X	icbtls			
isel	31 (0x1F)			rD					rA					rB					crb					0	1	1	1	1	0	X	isel			
mbar	31 (0x1F)			MO					///					1	1	0	1	0	1	0	1	1	0	1	1	0	/	X	mbar					
mfdcr	31 (0x1F)			rD					DCRN5-9					DCRN0-4					0	1	0	1	0	0	0	0	1	1	/	AFX	mfdcr			
mfpmr	31 (0x1F)			rD					PMRN5-9					PMRN0-4					0	1	0	1	0	0	1	1	1	0	0	AFX	mfpmr			
msync	31 (0x1F)			///					///					1	0	0	1	0	1	0	1	1	0	/	X	msync								
mtdcr	31 (0x1F)			rS					DCRN5-9					DCRN0-4					0	1	1	1	0	0	0	0	1	1	/	AFX	mtdcr			
mtpmr	31 (0x1F)			rS					PMRN5-9					PMRN0-4					0	1	1	1	0	0	1	1	1	0	0	AFX	mtpmr			
tlbivax	31 (0x1F)			///					rA					rB					1	1	0	0	0	1	0	0	1	0	/	X	tlbivax			
tlbre	31 (0x1F)			/// <sup>(2)</sup>					/// <sup>(2)</sup>					1	1	1	0	1	1	0	0	1	0	/	X	tlbre								
tlbsx	31 (0x1F)			/// <sup>(5)</sup>					rA					rB					1	1	1	0	0	1	0	0	1	0	/ <sup>5</sup>	X	tlbsx			
tlbwe	31 (0x1F)			/// <sup>(6)</sup>					/// <sup>(6)</sup>					1	1	1	1	0	1	0	0	1	0	/	X	tlbwe								
wait	31 (0x1F)			///					///					0	0	0	0	1	1	1	1	1	0	/		wait								
wrtree	31 (0x1F)			rS					///					0	0	1	0	0	0	0	0	0	1	1	/	X	wrtree							
wrtreei	31 (0x1F)			///					E	///					0	0	1	0	1	0	0	0	1	1	/	X	wrtreei							
cmp	31 (0x1F)			crfD					/	L	rA					rB					0	0	0	0	0	0	0	0	0	0	/	X	cmp	
tw	31 (0x1F)			TO					rA					rB					0	0	0	0	0	0	0	1	0	0	/	X	tw			
subfc	31 (0x1F)			rD					rA					rB					0	0	0	0	0	0	1	0	0	0	0	X	subfc			
subfc.	31 (0x1F)			rD					rA					rB					0	0	0	0	0	0	1	0	0	0	1	X	subfc.			





Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
addc	31 (0x1F)			rD					rA					rB					0	0	0	0	0	0	1	0	1	0	0	X	addc			
addc.	31 (0x1F)			rD					rA					rB					0	0	0	0	0	0	1	0	1	0	1	X	addc.			
mulhwu	31 (0x1F)			rD					rA					rB					/	0	0	0	0	0	1	0	1	1	0	X	mulhwu			
mulhwu.	31 (0x1F)			rD					rA					rB					/	0	0	0	0	0	1	0	1	1	1	X	mulhwu.			
mfcr	31 (0x1F)			rD					///										0	0	0	0	0	1	0	0	1	1	/	X	mfcr			
lwarx	31 (0x1F)			rD					rA					rB					0	0	0	0	0	1	0	1	0	0	/	X	lwarx			
lwzx	31 (0x1F)			rD					rA					rB					0	0	0	0	0	1	0	1	1	1	/	X	lwzx			
slw	31 (0x1F)			rS					rA					rB					0	0	0	0	0	1	1	0	0	0	0	X	slw			
slw.	31 (0x1F)			rS					rA					rB					0	0	0	0	0	1	1	0	0	0	1	X	slw.			
cntlzw	31 (0x1F)			rS					rA					///					0	0	0	0	0	1	1	0	1	0	0	X	cntlzw			
cntlzw.	31 (0x1F)			rS					rA					///					0	0	0	0	0	1	1	0	1	0	1	X	cntlzw.			
and	31 (0x1F)			rS					rA					rB					0	0	0	0	0	1	1	1	0	0	0	X	and			
and.	31 (0x1F)			rS					rA					rB					0	0	0	0	0	1	1	1	0	0	1	X	and.			
cmpl	31 (0x1F)			/	L	rA					rB					///					0	0	0	0	1	0	0	0	0	0	/	X	cmpl	
subf	31 (0x1F)			rD					rA					rB					0	0	0	0	1	0	1	0	0	0	0	X	subf			
subf.	31 (0x1F)			rD					rA					rB					0	0	0	0	1	0	1	0	0	0	1	X	subf.			
dcbst	31 (0x1F)			///					rA					rB					0	0	0	0	1	1	0	1	1	0	/	X	dcbst			
lwzux	31 (0x1F)			rD					rA					rB					0	0	0	0	1	1	0	1	1	1	/	X	lwzux			
andc	31 (0x1F)			rS					rA					rB					0	0	0	0	1	1	1	1	0	0	0	X	andc			
andc.	31 (0x1F)			rS					rA					rB					0	0	0	0	1	1	1	1	0	0	1	X	andc.			
mulhw	31 (0x1F)			rD					rA					rB					/	0	0	1	0	0	1	0	1	1	0	X	mulhw			
mulhw.	31 (0x1F)			rD					rA					rB					/	0	0	1	0	0	1	0	1	1	1	X	mulhw.			
mfmsr <sup>(1)</sup>	31 (0x1F)			rD					///										0	0	0	1	0	1	0	0	1	1	/	X	mfmsr			
dcbf	31 (0x1F)			///					rA					rB					0	0	0	1	0	1	0	1	1	0	/	X	dcbf			
lbzx	31 (0x1F)			rD					rA					rB					0	0	0	1	0	1	0	1	1	1	/	X	lbzx			
neg	31 (0x1F)			rD					rA					///					0	0	0	1	1	0	1	0	0	0	0	X	neg			
neg.	31 (0x1F)			rD					rA					///					0	0	0	1	1	0	1	0	0	0	1	X	neg.			
lbzux	31 (0x1F)			rD					rA					rB					0	0	0	1	1	1	0	1	1	1	/	X	lbzux			



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
nor			31 (0x1F)						rS					rA					rB			0	0	0	1	1	1	1	1	0	0	0	X	nor
nor.			31 (0x1F)						rS					rA					rB			0	0	0	1	1	1	1	1	0	0	1	X	nor.
subfe			31 (0x1F)						rD					rA					rB			0	0	1	0	0	0	1	0	0	0	0	X	subfe
subfe.			31 (0x1F)						rD					rA					rB			0	0	1	0	0	0	1	0	0	0	1	X	subfe.
adde			31 (0x1F)						rD					rA					rB			0	0	1	0	0	0	1	0	1	0	0	X	adde
adde.			31 (0x1F)						rD					rA					rB			0	0	1	0	0	0	1	0	1	0	1	X	adde.
mtrcf			31 (0x1F)						rS	/				CRM					/			0	0	1	0	0	1	0	0	0	0	/	XFX	mtrcf
mtmsr <sup>(1)</sup>			31 (0x1F)						rS					///								0	0	1	0	0	1	0	0	1	0	/	X	mtmsr
stwcx.			31 (0x1F)						rS					rA					rB			0	0	1	0	0	1	0	1	1	0	1	X	stwcx.
stwx			31 (0x1F)						rS					rA					rB			0	0	1	0	0	1	0	1	1	1	/	D	stwx
stwux			31 (0x1F)						rS					rA					rB			0	0	1	0	1	1	0	1	1	1	/	D	stwux
subfze			31 (0x1F)						rD					rA					///			0	0	1	1	0	0	1	0	0	0	0	X	subfze
subfze.			31 (0x1F)						rD					rA					///			0	0	1	1	0	0	1	0	0	0	1	X	subfze.
addze			31 (0x1F)						rD					rA					///			0	0	1	1	0	0	1	0	1	0	0	X	addze
addze.			31 (0x1F)						rD					rA					///			0	0	1	1	0	0	1	0	1	0	1	X	addze.
stbx			31 (0x1F)						rS					rA					rB			0	0	1	1	0	1	0	1	1	1	0	X	stbx
subfme			31 (0x1F)						rD					rA					///			0	0	1	1	1	0	1	0	0	0	0	X	subfme
subfme.			31 (0x1F)						rD					rA					///			0	0	1	1	1	0	1	0	0	0	1	X	subfme.
addme			31 (0x1F)						rD					rA					///			0	0	1	1	1	0	1	0	1	0	0	X	addme
addme.			31 (0x1F)						rD					rA					///			0	0	1	1	1	0	1	0	1	0	1	X	addme.
mullw			31 (0x1F)						rD					rA					rB			0	0	1	1	1	0	1	0	1	1	0	X	mullw
mullw.			31 (0x1F)						rD					rA					rB			0	0	1	1	1	0	1	0	1	1	1	X	mullw.
dcbtst			31 (0x1F)						CT					rA					rB			0	0	1	1	1	1	0	1	1	0	/	X	dcbtst
stbux			31 (0x1F)						rS					rA					rB			0	0	1	1	1	1	0	1	1	1	0	X	stbux
add			31 (0x1F)						rD					rA					rB			0	1	0	0	0	0	1	0	1	0	0	X	add
add.			31 (0x1F)						rD					rA					rB			0	1	0	0	0	0	1	0	1	0	1	X	add.
dcbt			31 (0x1F)						CT					rA					rB			0	1	0	0	0	1	0	1	1	0	/	X	dcbt
lhzx			31 (0x1F)						rD					rA					rB			0	1	0	0	0	1	0	1	1	1	/	X	lhzx



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
eqv			31 (0x1F)						rD					rA					rB			0	1	0	0	0	1	1	1	0	0	0	X	eqv
eqv.			31 (0x1F)						rD					rA					rB			0	1	0	0	0	1	1	1	0	0	1	X	eqv.
tlbie <sup>(1), (2)</sup>			31 (0x1F)						///					///					rB			0	1	0	0	1	1	0	0	1	0	0	X	tlbie
lhzux			31 (0x1F)						rD					rA					rB			0	1	0	0	1	1	0	1	1	1	/	X	lhzux
xor			31 (0x1F)						rS					rA					rB			0	1	0	0	1	1	1	1	0	0	0	X	xor
xor.			31 (0x1F)						rS					rA					rB			0	1	0	0	1	1	1	1	0	0	1	X	xor.
mfspr <sup>(7)</sup>			31 (0x1F)						rD					SPR[5–9]					SPR[0–4]			0	1	0	1	0	1	0	0	1	1	/	XF	mfspr
lhax			31 (0x1F)						rD					rA					rB			0	1	0	1	0	1	0	1	1	1	/	X	lhax
lhaux			31 (0x1F)						rD					rA					rB			0	1	0	1	1	1	0	1	1	1	/	X	lhaux
sthx			31 (0x1F)						rS					rA					rB			0	1	1	0	0	1	0	1	1	1	/	X	sthx
orc			31 (0x1F)						rS					rA					rB			0	1	1	0	0	1	1	1	0	0	0	X	orc
orc.			31 (0x1F)						rS					rA					rB			0	1	1	0	0	1	1	1	0	0	1	X	orc.
sthux			31 (0x1F)						rS					rA					rB			0	1	1	0	1	1	0	1	1	1	/	X	sthux
or			31 (0x1F)						rS					rA					rB			0	1	1	0	1	1	1	1	0	0	0	X	or
or.			31 (0x1F)						rS					rA					rB			0	1	1	0	1	1	1	1	0	0	1	X	or.
divwu			31 (0x1F)						rD					rA					rB			0	1	1	1	0	0	1	0	1	1	0	X	divwu
divwu.			31 (0x1F)						rD					rA					rB			0	1	1	1	0	0	1	0	1	1	1	X	divwu.
mtspr <sup>(2)</sup>			31 (0x1F)						rS					SPR[5–9]					SPR[0–4]			0	1	1	1	0	1	0	0	1	1	/	XF	mtspr
dcbi <sup>(1)</sup>			31 (0x1F)						///					rA					rB			0	1	1	1	0	1	0	1	1	0	/	X	dcbi
nand			31 (0x1F)						rS					rA					rB			0	1	1	1	0	1	1	1	0	0	0	X	nand
nand.			31 (0x1F)						rS					rA					rB			0	1	1	1	0	1	1	1	0	0	1	X	nand.
divw			31 (0x1F)						rD					rA					rB			0	1	1	1	1	0	1	0	1	1	0	X	divw
divw.			31 (0x1F)						rD					rA					rB			0	1	1	1	1	0	1	0	1	1	1	X	divw.
mcrxr			31 (0x1F)					crfD						///								1	0	0	0	0	0	0	0	0	0	/	X	mcrxr
subfco			31 (0x1F)						rD					rA					rB			1	0	0	0	0	0	1	0	0	0	0	X	subfco
subfco.			31 (0x1F)						rD					rA					rB			1	0	0	0	0	0	1	0	0	0	1	X	subfco.
addco			31 (0x1F)						rD					rA					rB			1	0	0	0	0	0	1	0	1	0	0	X	addco
addco.			31 (0x1F)						rD					rA					rB			1	0	0	0	0	0	1	0	1	0	1	X	addco.



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
lwbrx			31 (0x1F)						rD				rA						rB			1	0	0	0	0	1	0	1	1	0	/	X	lwbrx
srw			31 (0x1F)						rS				rA						rB			1	0	0	0	0	1	1	0	0	0	0	X	srw
srw.			31 (0x1F)						rS				rA						rB			1	0	0	0	0	1	1	0	0	0	1	X	srw.
subfo			31 (0x1F)						rD				rA						rB			1	0	0	0	1	0	1	0	0	0	0	X	subfo
subfo.			31 (0x1F)						rD				rA						rB			1	0	0	0	1	0	1	0	0	0	1	X	subfo.
tlbsync <sup>(1),(6)</sup>			31 (0x1F)						///				///						///			1	0	0	0	1	1	0	1	1	0	/	X	tlbsync
nego MBC			31 (0x1F)						rD				rA						///			1	0	0	1	1	0	1	0	0	0	0	X	nego
nego.			31 (0x1F)						rD				rA						///			1	0	0	1	1	0	1	0	0	0	1	X	nego.
subfeo			31 (0x1F)						rD				rA						rB			1	0	1	0	0	0	1	0	0	0	0	X	subfeo
subfeo.			31 (0x1F)						rD				rA						rB			1	0	1	0	0	0	1	0	0	0	1	X	subfeo.
addeo			31 (0x1F)						rD				rA						rB			1	0	1	0	0	0	1	0	1	0	0	X	addeo
addeo.			31 (0x1F)						rD				rA						rB			1	0	1	0	0	0	1	0	1	0	1	X	addeo.
stwbrx			31 (0x1F)						rS				rA						rB			1	0	1	0	0	1	0	1	1	0	/	X	stwbrx
subfzeo			31 (0x1F)						rD				rA						///			1	0	1	1	0	0	1	0	0	0	0	X	subfzeo
subfzeo.			31 (0x1F)						rD				rA						///			1	0	1	1	0	0	1	0	0	0	1	X	subfzeo.
addzeo			31 (0x1F)						rD				rA						///			1	0	1	1	0	0	1	0	1	0	0	X	addzeo
addzeo.			31 (0x1F)						rD				rA						///			1	0	1	1	0	0	1	0	1	0	1	X	addzeo.
subfmeo			31 (0x1F)						rD				rA						///			1	0	1	1	1	0	1	0	0	0	0	X	subfmeo
subfmeo.			31 (0x1F)						rD				rA						///			1	0	1	1	1	0	1	0	0	0	1	X	subfmeo.
addmeo			31 (0x1F)						rD				rA						///			1	0	1	1	1	0	1	0	1	0	0	X	addmeo
addmeo.			31 (0x1F)						rD				rA						///			1	0	1	1	1	0	1	0	1	0	1	X	addmeo.
mullwo			31 (0x1F)						rD				rA						rB			1	0	1	1	1	0	1	0	1	1	0	X	mullwo
mullwo.			31 (0x1F)						rD				rA						rB			1	0	1	1	1	0	1	0	1	1	1	X	mullwo.
dcba <sup>(6)</sup>			31 (0x1F)						///				rA						rB			1	0	1	1	1	1	0	1	1	0	/	X	dcba
addo			31 (0x1F)						rD				rA						rB			1	1	0	0	0	0	1	0	1	0	0	X	addo
addo.			31 (0x1F)						rD				rA						rB			1	1	0	0	0	0	1	0	1	0	1	X	addo.
lhbrx			31 (0x1F)						rD				rA						rB			1	1	0	0	0	1	0	1	1	0	/	X	lhbrx
sraw			31 (0x1F)						rS				rA						rB			1	1	0	0	0	1	1	0	0	0	X	sraw	



Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
sraw.			31 (0x1F)						rS					rA					rB			1	1	0	0	0	1	1	0	0	0	1	X	sraw.	
srawi			31 (0x1F)						rS					rA					SH			1	1	0	0	1	1	1	0	0	0	0	X	srawi	
srawi.			31 (0x1F)						rS					rA					SH			1	1	0	0	1	1	1	0	0	0	1	X	srawi.	
sthbrx			31 (0x1F)						rS					rA					rB			1	1	1	0	0	1	0	1	1	0	/	X	sthbrx	
extsh			31 (0x1F)						rS					rA					///			1	1	1	0	0	1	1	0	1	0	0	X	extsh	
extsh.			31 (0x1F)						rS					rA					///			1	1	1	0	0	1	1	0	1	0	1	X	extsh.	
extsb			31 (0x1F)						rS					rA					///			1	1	1	0	1	1	1	0	1	0	0	X	extsb	
extsb.			31 (0x1F)						rS					rA					///			1	1	1	0	1	1	1	0	1	0	1	X	extsb.	
divwuo			31 (0x1F)						rD					rA					rB			1	1	1	1	0	0	1	0	1	1	0	X	divwuo	
divwuo.			31 (0x1F)						rD					rA					rB			1	1	1	1	0	0	1	0	1	1	1	X	divwuo.	
icbi			31 (0x1F)						///					rA					rB			1	1	1	1	0	1	0	1	1	0	/	X	icbi	
divwo			31 (0x1F)						rD					rA					rB			1	1	1	1	1	0	1	0	1	1	0	X	divwo	
divwo.			31 (0x1F)						rD					rA					rB			1	1	1	1	1	0	1	0	1	1	1	X	divwo.	
dcbz			31 (0x1F)						///					rA					rB			1	1	1	1	1	1	0	1	1	0	/	X	dcbz	
lwz			32 (0x20)						rD					rA																			D	lwz	
lwzu			33 (0x21)						rD					rA																				D	lwzu
lbz			34(0x22)						rD					rA																				D	lbz
lbzu			35(0x23)						rD					rA																				D	lbzu
stw			36(0x24)						rS					rA																				D	stw
stwu			37(0x25)						rS					rA																				D	stwu
stb			38(0x26)						rS					rA																				D	stb
stbu			39(0x27)						rS					rA																				D	stbu
lhz			40(0x28)						rD					rA																				D	lhz
lhzu			41(0x29)						rD					rA																				D	lhzu
lha			42(0x2A)						rD					rA																				D	lha
lhau			43(0x2B)						rD					rA																				D	lhau
sth			44(0x2C)						rS					rA																				D	sth
sthu			45(0x2D)						rS					rA																				D	sthu
lmw			46(0x2E)						rD					rA																				D	lmw


**Table 270. Instructions sorted by primary opcodes (decimal and hexadecimal) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
stmw	47(0x2F)				rS				rA				D																D	stmw				
fres <sup>(6)</sup>	59(0x3B)				frD				///				frB				///				1	1	0	0	0	0	1	A	fres					
fres. <sup>(6)</sup>	59(0x3B)				frD				///				frB				///				1	1	0	0	0	1	A	fres.						
fsel <sup>(6)</sup>	63(0x3F)				frD				frA				frB				frC				1	0	1	1	1	0	A	fsel						
fsel. <sup>(6)</sup>	63(0x3F)				frD				frA				frB				frC				1	0	1	1	1	1	A	fsel.						

1. Supervisor-level instruction.
2. d = UIMM \* 8.
3. d = UIMM \* 2.
4. d = UIMM \* 4.
5. This field is defined as allocated by the Book E architecture, for possible use in an implementation.
6. Optional to the PowerPC classic architecture.
7. Access level is determined by whether the SPR is defined as a user- or supervisor-level SPR.

### A.3 Instructions sorted by mnemonic (binary)

[Table 271](#) lists instructions in alphabetical order by mnemonic with binary values. This list also includes simplified mnemonics and their equivalents using standard mnemonics.



Table 271. Instructions sorted by mnemonic (binary)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
add	0	1	1	1	1	1			rD			rA							rB		0	1	0	0	0	0	0	1	0	1	0	0	X	add
add.	0	1	1	1	1	1			rD			rA							rB		0	1	0	0	0	0	0	1	0	1	0	1	X	add.
addc	0	1	1	1	1	1			rD			rA							rB		0	0	0	0	0	0	0	1	0	1	0	0	X	addc
addc.	0	1	1	1	1	1			rD			rA							rB		0	0	0	0	0	0	0	1	0	1	0	1	X	addc.
addco	0	1	1	1	1	1			rD			rA							rB		1	0	0	0	0	0	0	1	0	1	0	0	X	addco
addco.	0	1	1	1	1	1			rD			rA							rB		1	0	0	0	0	0	0	1	0	1	0	1	X	addco.
adde	0	1	1	1	1	1			rD			rA							rB		0	0	1	0	0	0	0	1	0	1	0	0	X	adde
adde.	0	1	1	1	1	1			rD			rA							rB		0	0	1	0	0	0	0	1	0	1	0	1	X	adde.
addeo	0	1	1	1	1	1			rD			rA							rB		1	0	1	0	0	0	0	1	0	1	0	0	X	addeo
addeo.	0	1	1	1	1	1			rD			rA							rB		1	0	1	0	0	0	0	1	0	1	0	1	X	addeo.
addi	0	0	1	1	1	0			rD			rA		SIMM													D	addi						
addic	0	0	1	1	0	0			rD			rA		SIMM													D	addic						
addic.	0	0	1	1	0	1			rD			rA		SIMM													D	addic.						
addis	0	0	1	1	1	1			rD			rA		SIMM													D	addis						
addme	0	1	1	1	1	1			rD			rA		///							0	0	1	1	1	0	1	0	1	0	0	X	addme	
addme.	0	1	1	1	1	1			rD			rA		///							0	0	1	1	1	0	1	0	1	0	1	X	addme.	
addmeo	0	1	1	1	1	1			rD			rA		///							1	0	1	1	1	0	1	0	1	0	0	X	addmeo	
addmeo.	0	1	1	1	1	1			rD			rA		///							1	0	1	1	1	0	1	0	1	0	1	X	addmeo.	
addo	0	1	1	1	1	1			rD			rA							rB		1	1	0	0	0	0	0	1	0	1	0	0	X	addo
addo.	0	1	1	1	1	1			rD			rA							rB		1	1	0	0	0	0	0	1	0	1	0	1	X	addo.
addze	0	1	1	1	1	1			rD			rA		///							0	0	1	1	0	0	1	0	1	0	0	X	addze	
addze.	0	1	1	1	1	1			rD			rA		///							0	0	1	1	0	0	1	0	1	0	1	X	addze.	
addzeo	0	1	1	1	1	1			rD			rA		///							1	0	1	1	0	0	1	0	1	0	0	X	addzeo	
addzeo.	0	1	1	1	1	1			rD			rA		///							1	0	1	1	0	0	1	0	1	0	1	X	addzeo.	
and	0	1	1	1	1	1			rS			rA							rB		0	0	0	0	0	0	1	1	1	0	0	0	X	and
and.	0	1	1	1	1	1			rS			rA							rB		0	0	0	0	0	0	1	1	1	0	0	1	X	and.
andc	0	1	1	1	1	1			rS			rA							rB		0	0	0	0	1	1	1	1	0	0	0	X	andc	
andc.	0	1	1	1	1	1			rS			rA							rB		0	0	0	0	1	1	1	1	0	0	1	X	andc.	





Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
andi.	0	1	1	1	0	0	rS			rA			UIMM															D	andi.					
andis.	0	1	1	1	0	1	rS			rA			UIMM															D	andis.					
b	0	1	0	0	1	0	LI															0	0	I	b									
ba	0	1	0	0	1	0	LI															1	0	I	ba									
bc	0	1	0	0	0	0	BO			BI			BD															0	0	B	bc			
bca	0	1	0	0	0	0	BO			BI			BD															1	0	B	bca			
bcctr	0	1	0	0	1	1	BO			BI			///			1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	XL	bcctr	
bcctrl	0	1	0	0	1	1	BO			BI			///			1	0	0	0	0	0	1	0	0	0	0	0	0	0	1	XL	bcctrl		
bcl	0	1	0	0	0	0	BO			BI			BD															0	1	B	bcl			
bcla	0	1	0	0	0	0	BO			BI			BD															1	1	B	bcla			
bclr	0	1	0	0	1	1	BO			BI			///			0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	XL	bclr	
bclrl	0	1	0	0	1	1	BO			BI			///			0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	XL	bclrl		
bctr	<b>bctr</b> <sup>(1)</sup>						equivalent to			<b>bcctr 20,0</b>																	bctr							
bctrl	<b>bctrl</b> <sup>(1)</sup>						equivalent to			<b>bcctrl 20,0</b>																	bctrl							
bdnz	<b>bdnz target</b> <sup>(1)</sup>						equivalent to			<b>bc 16,0,target</b>																	bdnz							
bdnza	<b>bdnza target</b> <sup>(1)</sup>						equivalent to			<b>bca 16,0,target</b>																	bdnza							
bdnzf	<b>bdnzf BI,target</b>						equivalent to			<b>bc 0,BI,target</b>																	bdnzf							
bdnzfa	<b>bdnzfa BI,target</b>						equivalent to			<b>bca 0,BI,target</b>																	bdnzfa							
bdnzfl	<b>bdnzfl BI,target</b>						equivalent to			<b>bcl 0,BI,target</b>																	bdnzfl							
bdnzfla	<b>bdnzfla BI,target</b>						equivalent to			<b>bcla 0,BI,target</b>																	bdnzfla							
bdnzflr	<b>bdnzflr BI</b>						equivalent to			<b>bclr 0,BI</b>																	bdnzflr							
bdnzflrl	<b>bdnzflrl BI</b>						equivalent to			<b>bclrl 0,BI</b>																	bdnzflrl							
bdnzl	<b>bdnzl target</b> <sup>(1)</sup>						equivalent to			<b>bcl 16,0,target</b>																	bdnzl							
bdnzla	<b>bdnzla target</b> <sup>(1)</sup>						equivalent to			<b>bcla 16,0,target</b>																	bdnzla							
bdnzlr	<b>bdnzlr BI</b>						equivalent to			<b>bclr 16,BI</b>																	bdnzlr							
bdnzlrl	<b>bdnzlrl</b> <sup>(1)</sup>						equivalent to			<b>bclrl 16,0</b>																	bdnzlrl							
bdnzt	<b>bdnzt BI,target</b>						equivalent to			<b>bc 8,BI,target</b>																	bdnzt							
bdnzta	<b>bdnzta BI,target</b>						equivalent to			<b>bca 8,BI,target</b>																	bdnzta							
bdnztl	<b>bdnztl BI,target</b>						equivalent to			<b>bcl 8,0,target</b>																	bdnztl							



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic										
<b>bdnztla</b>	<b>bdnztla BI,target</b> equivalent to										<b>bcla 8,BI,target</b>																												<b>bdnztla</b>					
<b>bdnztlr</b>	<b>bdnztlr BI</b> equivalent to										<b>bclr 8,BI</b>																														<b>bdnztlr</b>			
<b>bdnztlr</b>	<b>bdnztlr BI</b> equivalent to										<b>bclr 8,BI</b>																														<b>bdnztlr</b>			
<b>bdnztlrl</b>	<b>bdnztlrl BI</b> equivalent to										<b>bclrl 8,BI</b>																														<b>bdnztlrl</b>			
<b>bdz</b>	<b>bdz target</b> <sup>(1)</sup> equivalent to										<b>bc 18,0,target</b>																														<b>bdz</b>			
<b>bdza</b>	<b>bdza target</b> <sup>(1)</sup> equivalent to										<b>bca 18,0,target</b>																															<b>bdza</b>		
<b>bdzf</b>	<b>bdzf BI,target</b> equivalent to										<b>bc 2,BI,target</b>																															<b>bdzf</b>		
<b>bdzfa</b>	<b>bdzfa BI,target</b> equivalent to										<b>bca 2,BI,target</b>																																<b>bdzfa</b>	
<b>bdzfl</b>	<b>bdzfl BI,target</b> equivalent to										<b>bcl 2,BI,target</b>																															<b>bdzfl</b>		
<b>bdzfla</b>	<b>bdzfla BI,target</b> equivalent to										<b>bcla 2,BI,target</b>																															<b>bdzfla</b>		
<b>bdzflr</b>	<b>bdzflr BI</b> equivalent to										<b>bclr 2,BI</b>																															<b>bdzflr</b>		
<b>bdzflrl</b>	<b>bdzflrl BI</b> equivalent to										<b>bclrl 2,BI</b>																															<b>bdzflrl</b>		
<b>bdzl</b>	<b>bdzl target</b> <sup>(1)</sup> equivalent to										<b>bcl 18,BI,target</b>																															<b>bdzl</b>		
<b>bdzla</b>	<b>bdzla target</b> <sup>(1)</sup> equivalent to										<b>bcla 18,BI,target</b>																																<b>bdzla</b>	
<b>bdzlr</b>	<b>bdzlr</b> <sup>(1)</sup> equivalent to										<b>bclr 18,0</b>																															<b>bdzlr</b>		
<b>bdzlr</b>	<b>bdzlr</b> <sup>(1)</sup> equivalent to										<b>bclrl 18,0</b>																															<b>bdzlr</b>		
<b>bdzt</b>	<b>bdzt BI,target</b> equivalent to										<b>bc 10,BI,target</b>																																<b>bdzt</b>	
<b>bdzta</b>	<b>bdzta BI,target</b> equivalent to										<b>bca 10,BI,target</b>																																	<b>bdzta</b>
<b>bdztl</b>	<b>bdztl BI,target</b> equivalent to										<b>bcl 10,BI,target</b>																																<b>bdztl</b>	
<b>bdztle</b>	<b>bdztle BI,target</b> equivalent to										<b>bcla 10,BI,target</b>																																<b>bdztle</b>	
<b>bdztlrl</b>	<b>bdztlrl BI</b> equivalent to										<b>bclrl 10, BI</b>																																<b>bdztlrl</b>	
<b>beq</b>	<b>beq crS,target</b> equivalent to										<b>bc 12,BI<sup>(2)</sup>,target</b>																															<b>beq</b>		
<b>beqa</b>	<b>beqa crS,target</b> equivalent to										<b>bca 12,BI<sup>(2)</sup>,target</b>																																<b>beqa</b>	
<b>beqctr</b>	<b>beqctr crS,target</b> equivalent to										<b>bcctr 12,BI<sup>(2)</sup>,target</b>																																<b>beqctr</b>	
<b>beqctrl</b>	<b>beqctrl crS,target</b> equivalent to										<b>bcctrl 12,BI<sup>(2)</sup>,target</b>																																<b>beqctrl</b>	
<b>beql</b>	<b>beql crS,target</b> equivalent to										<b>bcl 12,BI<sup>(2)</sup>,target</b>																															<b>beql</b>		
<b>beqla</b>	<b>beqla crS,target</b> equivalent to										<b>bcla 12,BI<sup>(2)</sup>,target</b>																																<b>beqla</b>	
<b>beqlr</b>	<b>beqlr crS,target</b> equivalent to										<b>bclr 12,BI<sup>(2)</sup>,target</b>																															<b>beqlr</b>		
<b>beqlrl</b>	<b>beqlrl crS,target</b> equivalent to										<b>bclrl 12,BI<sup>(2)</sup>,target</b>																															<b>beqlrl</b>		



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
bf	bf BI,target		equivalent to				bc 4,BI,target																bf											
bfa	bfa BI,target		equivalent to				bca 4,BI,target																bfa											
bfctr	bfctr BI		equivalent to				bcctr 4,BI																bfctr											
bfctrl	bfctrl BI		equivalent to				bcctrl 4,BI																bfctrl											
bfl	bfl BI,target		equivalent to				bcl 4,BI,target																bfl											
bfla	bfla BI,target		equivalent to				bcla 4,BI,target																bfla											
bflr	bflr BI		equivalent to				bclr 4,BI																bflr											
bflrl	bflrl BI		equivalent to				bcrlr 4,BI																bflrl											
bge	bge crS,target		equivalent to				bc 4,BI <sup>(3)</sup> ,target																bge											
bgea	bgea crS,target		equivalent to				bca 4,BI <sup>(3)</sup> ,target																bgea											
bgectr	bgectr crS,target		equivalent to				bcctr 4,BI <sup>(3)</sup> ,target																bgectr											
bgectrl	bgectrl crS,target		equivalent to				bcctrl 4,BI <sup>(3)</sup> ,target																bgectrl											
bgel	bgel crS,target		equivalent to				bcl 4,BI <sup>(3)</sup> ,target																bgel											
bgela	bgela crS,target		equivalent to				bcla 4,BI <sup>(3)</sup> ,target																bgela											
bgelr	bgelr crS,target		equivalent to				bclr 4,BI <sup>(3)</sup> ,target																bgelr											
bgelrl	bgelrl crS,target		equivalent to				bcrlr 4,BI <sup>(3)</sup> ,target																bgelrl											
bgt	bgt crS,target		equivalent to				bc 12,BI <sup>(4)</sup> ,target																bgt											
bgt a	bgt a crS,target		equivalent to				bca 12,BI <sup>(4)</sup> ,target																bgt a											
bgtctr	bgtctr crS,target		equivalent to				bcctr 12,BI <sup>(4)</sup> ,target																bgtctr											
bgtctrl	bgtctrl crS,target		equivalent to				bcctrl 12,BI <sup>(4)</sup> ,target																bgtctrl											
bgtl	bgtl crS,target		equivalent to				bcl 12,BI <sup>(4)</sup> ,target																bgtl											
bgtla	bgtla crS,target		equivalent to				bcla 12,BI <sup>(4)</sup> ,target																bgtla											
bgtlr	bgtlr crS,target		equivalent to				bclr 12,BI <sup>(4)</sup> ,target																bgtlr											
bgtlrl	bgtlrl crS,target		equivalent to				bcrlr 12,BI <sup>(4)</sup> ,target																bgtlrl											
bl	0	1	0	0	1	0														LI	0	1	I	bl										
bla	0	1	0	0	1	0														LI	1	1	I	bla										
ble	ble crS,target		equivalent to				bc 4,BI <sup>(4)</sup> ,target																ble											
blea	blea crS,target		equivalent to				bca 4,BI <sup>(4)</sup> ,target																blea											
blectr	blectr crS,target		equivalent to				bcctr 4,BI <sup>(4)</sup> ,target																blectr											



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic									
<b>blectrl</b>	<b>blectrl crS,target</b> equivalent to										<b>bcctrl 4,BI<sup>(4)</sup>,target</b>																													<b>blectrl</b>			
<b>blel</b>	<b>blel crS,target</b> equivalent to										<b>bcl 4,BI<sup>(4)</sup>,target</b>																														<b>blel</b>		
<b>blela</b>	<b>blela crS,target</b> equivalent to										<b>bcla 4,BI<sup>(4)</sup>,target</b>																														<b>blela</b>		
<b>blelr</b>	<b>blelr crS,target</b> equivalent to										<b>bclr 4,BI<sup>(4)</sup>,target</b>																														<b>blelr</b>		
<b>blelrl</b>	<b>blelrl crS,target</b> equivalent to										<b>bclrl 4,BI<sup>(4)</sup>,target</b>																														<b>blelrl</b>		
<b>blr</b>	<b>blr<sup>(1)</sup></b> equivalent to										<b>bclr 20,0</b>																														<b>blr</b>		
<b>blrl</b>	<b>blrl<sup>(1)</sup></b> equivalent to										<b>bclrl 20,0</b>																															<b>blrl</b>	
<b>blt</b>	<b>blt crS,target</b> equivalent to										<b>bc 12,BI,target</b>																															<b>blt</b>	
<b>blta</b>	<b>blta crS,target</b> equivalent to										<b>bca 12,BI<sup>(3)</sup>,target</b>																															<b>blta</b>	
<b>bltctr</b>	<b>bltctr crS,target</b> equivalent to										<b>bcctr 12,BI<sup>(3)</sup>,target</b>																															<b>bltctr</b>	
<b>bltctrl</b>	<b>bltctrl crS,target</b> equivalent to										<b>bcctrl 12,BI<sup>(3)</sup>,target</b>																															<b>bltctrl</b>	
<b>bltl</b>	<b>bltl crS,target</b> equivalent to										<b>bcl 12,BI<sup>(3)</sup>,target</b>																															<b>bltl</b>	
<b>bltla</b>	<b>bltla crS,target</b> equivalent to										<b>bcla 12,BI<sup>(3)</sup>,target</b>																															<b>bltla</b>	
<b>bltlr</b>	<b>bltlr crS,target</b> equivalent to										<b>bclr 12,BI<sup>(3)</sup>,target</b>																															<b>bltlr</b>	
<b>bltlrl</b>	<b>bltlrl crS,target</b> equivalent to										<b>bclrl 12,BI<sup>(3)</sup>,target</b>																															<b>bltlrl</b>	
<b>bne</b>	<b>bne crS,target</b> equivalent to										<b>bc 4,BI<sup>(3)</sup>,target</b>																															<b>bne</b>	
<b>bnea</b>	<b>bnea crS,target</b> equivalent to										<b>bca 4,BI<sup>(3)</sup>,target</b>																																<b>bnea</b>
<b>bnctr</b>	<b>bnctr crS,target</b> equivalent to										<b>bcctr 4,BI<sup>(3)</sup>,target</b>																																<b>bnctr</b>
<b>bnctrl</b>	<b>bnctrl crS,target</b> equivalent to										<b>bcctrl 4,BI<sup>(3)</sup>,target</b>																																<b>bnctrl</b>
<b>bnel</b>	<b>bnel crS,target</b> equivalent to										<b>bcl 4,BI<sup>(3)</sup>,target</b>																																<b>bnel</b>
<b>bnela</b>	<b>bnela crS,target</b> equivalent to										<b>bcla 4,BI<sup>(3)</sup>,target</b>																																<b>bnela</b>
<b>bnelr</b>	<b>bnelr crS,target</b> equivalent to										<b>bclr 4,BI<sup>(3)</sup>,target</b>																																<b>bnelr</b>
<b>bnelrl</b>	<b>bnelrl crS,target</b> equivalent to										<b>bclrl 4,BI<sup>(3)</sup>,target</b>																																<b>bnelrl</b>
<b>bng</b>	<b>bng crS,target</b> equivalent to										<b>bc 4,BI<sup>(4)</sup>,target</b>																																<b>bng</b>
<b>bnga</b>	<b>bnga crS,target</b> equivalent to										<b>bca 4,BI<sup>(4)</sup>,target</b>																																<b>bnga</b>
<b>bngctr</b>	<b>bngctr crS,target</b> equivalent to										<b>bcctr 4,BI<sup>(4)</sup>,target</b>																																<b>bngctr</b>
<b>bngctrl</b>	<b>bngctrl crS,target</b> equivalent to										<b>bcctrl 4,BI<sup>(4)</sup>,target</b>																																<b>bngctrl</b>
<b>bngl</b>	<b>bngl crS,target</b> equivalent to										<b>bcl 4,BI<sup>(4)</sup>,target</b>																																<b>bngl</b>
<b>bngla</b>	<b>bngla crS,target</b> equivalent to										<b>bcla 4,BI<sup>(4)</sup>,target</b>																																<b>bngla</b>



**Table 271. Instructions sorted by mnemonic (binary) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic																		
<b>bnglr</b>	<b>bnglr crS,target</b> equivalent to																<b>bclr 4,BI<sup>(4)</sup>,target</b>																													<b>bnglr</b>						
<b>bnglrl</b>	<b>bnglrl crS,target</b> equivalent to																<b>bclrl 4,BI<sup>(4)</sup>,target</b>																														<b>bnglrl</b>					
<b>bnl</b>	<b>bnl crS,target</b> equivalent to																<b>bc 4,BI<sup>(3)</sup>,target</b>																														<b>bnl</b>					
<b>bnla</b>	<b>bnla crS,target</b> equivalent to																<b>bca 4,BI<sup>(3)</sup>,target</b>																															<b>bnla</b>				
<b>bnlctr</b>	<b>bnlctr crS,target</b> equivalent to																<b>bcctr 4,BI<sup>(3)</sup>,target</b>																															<b>bnlctr</b>				
<b>bnlctrl</b>	<b>bnlctrl crS,target</b> equivalent to																<b>bcctrl 4,BI<sup>(3)</sup>,target</b>																															<b>bnlctrl</b>				
<b>bnll</b>	<b>bnll crS,target</b> equivalent to																<b>bcl 4,BI<sup>(3)</sup>,target</b>																															<b>bnll</b>				
<b>bnlla</b>	<b>bnlla crS,target</b> equivalent to																<b>bcla 4,BI<sup>(3)</sup>,target</b>																																<b>bnlla</b>			
<b>bnllr</b>	<b>bnllr crS,target</b> equivalent to																<b>bclr 4,BI<sup>(3)</sup>,target</b>																																<b>bnllr</b>			
<b>bnllrl</b>	<b>bnllrl crS,target</b> equivalent to																<b>bclrl 4,BI<sup>(3)</sup>,target</b>																																<b>bnllrl</b>			
<b>bns</b>	<b>bns crS,target</b> equivalent to																<b>bc 4,BI<sup>(5)</sup>,target</b>																																<b>bns</b>			
<b>bnsa</b>	<b>bnsa crS,target</b> equivalent to																<b>bca 4,BI<sup>(5)</sup>,target</b>																																	<b>bnsa</b>		
<b>bnsctr</b>	<b>bnsctr crS,target</b> equivalent to																<b>bcctr 4,BI<sup>(5)</sup>,target</b>																																	<b>bnsctr</b>		
<b>bnsctrl</b>	<b>bnsctrl crS,target</b> equivalent to																<b>bcctrl 4,BI<sup>(5)</sup>,target</b>																																	<b>bnsctrl</b>		
<b>bnsI</b>	<b>bnsI crS,target</b> equivalent to																<b>bcl 4,BI<sup>(5)</sup>,target</b>																																	<b>bnsI</b>		
<b>bnsIa</b>	<b>bnsIa crS,target</b> equivalent to																<b>bcla 4,BI<sup>(5)</sup>,target</b>																																		<b>bnsIa</b>	
<b>bnsIrl</b>	<b>bnsIrl crS,target</b> equivalent to																<b>bclrl 4,BI<sup>(5)</sup>,target</b>																																		<b>bnsIrl</b>	
<b>bnu</b>	<b>bnu crS,target</b> equivalent to																<b>bc 4,BI<sup>(5)</sup>,target</b>																																		<b>bnu</b>	
<b>bnua</b>	<b>bnua crS,target</b> equivalent to																<b>bca 4,BI<sup>(5)</sup>,target</b>																																		<b>bnua</b>	
<b>bnuctr</b>	<b>bnuctr crS,target</b> equivalent to																<b>bcctr 4,BI<sup>(5)</sup>,target</b>																																	<b>bnuctr</b>		
<b>bnuctrl</b>	<b>bnuctrl crS,target</b> equivalent to																<b>bcctrl 4,BI<sup>(5)</sup>,target</b>																																		<b>bnuctrl</b>	
<b>bnul</b>	<b>bnul crS,target</b> equivalent to																<b>bcl 4,BI<sup>(5)</sup>,target</b>																																		<b>bnul</b>	
<b>bnula</b>	<b>bnula crS,target</b> equivalent to																<b>bcla 4,BI<sup>(5)</sup>,target</b>																																			<b>bnula</b>
<b>bnulr</b>	<b>bnulr crS,target</b> equivalent to																<b>bclr 4,BI<sup>(5)</sup>,target</b>																																		<b>bnulr</b>	
<b>bnulrl</b>	<b>bnulrl crS,target</b> equivalent to																<b>bclrl 4,BI<sup>(5)</sup>,target</b>																																		<b>bnulrl</b>	
<b>brinc</b>	0	0	0	1	0	0	rD				rA				rB				0	1	0	0	0	0	0	0	0	1	1	1	1	1	1	1	EVX	<b>brinc</b>																
<b>bso</b>	<b>bso crS,target</b> equivalent to																<b>bc 12,BI<sup>(5)</sup>,target</b>																															<b>bso</b>				
<b>bsoa</b>	<b>bsoa crS,target</b> equivalent to																<b>bca 12,BI<sup>(5)</sup>,target</b>																																<b>bsoa</b>			



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
<b>bsoctr</b>	<b>bsoctr</b> crS,target equivalent to						<b>bcctr</b> 12,BI <sup>(5)</sup> ,target																<b>bsoctr</b>											
<b>bsoctrl</b>	<b>bsoctrl</b> crS,targetequivalent to						<b>bcctrl</b> 12,BI <sup>(5)</sup> ,target																<b>bsoctrl</b>											
<b>bsol</b>	<b>bsol</b> crS,target equivalent to						<b>bcl</b> 12,BI <sup>(5)</sup> ,target																<b>bsol</b>											
<b>bsola</b>	<b>bsola</b> crS,target equivalent to						<b>bcla</b> 12,BI <sup>(5)</sup> ,target																<b>bsola</b>											
<b>bsolr</b>	<b>bsolr</b> crS,target equivalent to						<b>bclr</b> 12,BI <sup>(5)</sup> ,target																<b>bsolr</b>											
<b>bsolrl</b>	<b>bsolrl</b> crS,target equivalent to						<b>bclrl</b> 12,BI <sup>(5)</sup> ,target																<b>bsolrl</b>											
<b>bt</b>	<b>bt</b> BI,target equivalent to						<b>bc</b> 12,BI,target																<b>bt</b>											
<b>bta</b>	<b>bta</b> BI,target equivalent to						<b>bca</b> 12,BI,target																<b>bta</b>											
<b>btctr</b>	<b>btctr</b> BI equivalent to						<b>bcctr</b> 12,BI																<b>btctr</b>											
<b>btctrl</b>	<b>btctrl</b> BI equivalent to						<b>bcctrl</b> 12,BI																<b>btctrl</b>											
<b>btl</b>	<b>btl</b> BI,target equivalent to						<b>bcl</b> 12,BI,target																<b>btl</b>											
<b>btla</b>	<b>btla</b> BI,target equivalent to						<b>bcla</b> 12,BI,target																<b>btla</b>											
<b>btlr</b>	<b>btlr</b> BI equivalent to						<b>bclr</b> 12,BI																<b>btlr</b>											
<b>btlrl</b>	<b>btlrl</b> BI equivalent to						<b>bclrl</b> 12,BI																<b>btlrl</b>											
<b>bun</b>	<b>bun</b> crS,target equivalent to						<b>bc</b> 12,BI <sup>(5)</sup> ,target																<b>bun</b>											
<b>buna</b>	<b>buna</b> crS,target equivalent to						<b>bca</b> 12,BI <sup>(5)</sup> ,target																<b>buna</b>											
<b>bunctr</b>	<b>bunctr</b> crS,target equivalent to						<b>bcctr</b> 12,BI <sup>(5)</sup> ,target																<b>bunctr</b>											
<b>bunctrl</b>	<b>bunctrl</b> crS,targetequivalent to						<b>bcctrl</b> 12,BI <sup>(5)</sup> ,target																<b>bunctrl</b>											
<b>bunl</b>	<b>bunl</b> crS,target equivalent to						<b>bcl</b> 12,BI <sup>(5)</sup> ,target																<b>bunl</b>											
<b>bunla</b>	<b>bunla</b> crS,target equivalent to						<b>bcla</b> 12,BI <sup>(5)</sup> ,target																<b>bunla</b>											
<b>bunlr</b>	<b>bunlr</b> crS,target equivalent to						<b>bclr</b> 12,BI <sup>(5)</sup> ,target																<b>bunlr</b>											
<b>bunlrl</b>	<b>bunlrl</b> crS,target equivalent to						<b>bclrl</b> 12,BI <sup>(5)</sup> ,target																<b>bunlrl</b>											
<b>clrlslwi</b>	<b>clrlslwi</b> rA,rS,b,n (n ≤ b ≤ 31) equivalent to						<b>rlwinm</b> rA,rS,n,b – n,31 – n																<b>clrlslwi</b>											
<b>clrlwi</b>	<b>clrlwi</b> rA,rS,n (n < 32) equivalent to						<b>rlwinm</b> rA,rS,0,n,31																<b>clrlwi</b>											
<b>clrrwi</b>	<b>clrrwi</b> rA,rS,n (n < 32) equivalent to						<b>rlwinm</b> rA,rS,0,0,31 – n																<b>clrrwi</b>											
<b>cmp</b>	0	1	1	1	1	1	crfD	/	L	rA			rB			0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	/	X	<b>cmp</b>	
<b>cmpi</b>	0	0	1	0	1	1	crfD	/	L	rA			SIMM															D	<b>cmpi</b>					
<b>cmpl</b>	0	1	1	1	1	1	/	L	rA			rB			///			0	0	0	0	1	0	0	0	0	0	0	0	/	X	<b>cmpl</b>		
<b>cmpli</b>	0	0	1	0	1	0	crfD	/	L	rA			UIMM															D	<b>cmpli</b>					



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic		
<b>cmplw</b>	<b>cmplw crD,rA,rB</b>						equivalent to						<b>cmpl crD,0,rA,rB</b>												<b>cmplw</b>											
<b>cmplwi</b>	<b>cmplwi crD,rA,UIMM</b>						equivalent to						<b>cmpli crD,0,rA,UIMM</b>												<b>cmplwi</b>											
<b>cmpw</b>	<b>cmpw crD,rA,rB</b>						equivalent to						<b>cmp crD,0,rA,rB</b>												<b>cmpw</b>											
<b>cmpwi</b>	<b>cmpwi crD,rA,SIMM</b>						equivalent to						<b>cmpi crD,0,rA,SIMM</b>												<b>cmpwi</b>											
<b>cntlzw</b>	0	1	1	1	1	1																												X	<b>cntlzw</b>	
<b>cntlzw.</b>	0	1	1	1	1	1																												X	<b>cntlzw.</b>	
<b>crand</b>	0	1	0	0	1	1																											/	XL	<b>crand</b>	
<b>crandc</b>	0	1	0	0	1	1																											/	XL	<b>crandc</b>	
<b>crclr</b>	<b>crclr bx</b>						equivalent to						<b>crxor bx,bx,bx</b>												<b>crclr</b>											
<b>creqv</b>	0	1	0	0	1	1																											/	XL	<b>creqv</b>	
<b>crmve</b>	<b>crmve bx,by</b>						equivalent to						<b>cror bx,by,by</b>												<b>crmve</b>											
<b>crnand</b>	0	1	0	0	1	1																											/	XL	<b>crnand</b>	
<b>crnor</b>	0	1	0	0	1	1																											/	XL	<b>crnor</b>	
<b>crnot</b>	<b>crnot bx,by</b>						equivalent to						<b>crnor bx,by,by</b>												<b>crnot</b>											
<b>cror</b>	0	1	0	0	1	1																											/	XL	<b>cror</b>	
<b>crorc</b>	0	1	0	0	1	1																											/	XL	<b>crorc</b>	
<b>crset</b>	<b>crset bx</b>						equivalent to						<b>creqv bx,bx,bx</b>												<b>crset</b>											
<b>crxor</b>	0	1	0	0	1	1																											/	XL	<b>crxor</b>	
<b>dcba<sup>(6)</sup></b>	0	1	1	1	1	1																											/	X	<b>dcba</b>	
<b>dcbf</b>	0	1	1	1	1	1																											/	X	<b>dcbf</b>	
<b>dcbi<sup>(7)</sup></b>	0	1	1	1	1	1																											/	X	<b>dcbi</b>	
<b>dcblic</b>	0	1	1	1	1	1																													X	<b>dcblic</b>
<b>dcbst</b>	0	1	1	1	1	1																											/	X	<b>dcbst</b>	
<b>dcbt</b>	0	1	1	1	1	1																											/	X	<b>dcbt</b>	
<b>dcbtlls</b>	0	1	1	1	1	1																													X	<b>dcbtlls</b>
<b>dcbtst</b>	0	1	1	1	1	1																											/	X	<b>dcbtst</b>	
<b>dcbtstlls</b>	0	1	1	1	1	1																													X	<b>dcbtstlls</b>
<b>dcbz</b>	0	1	1	1	1	1																											/	X	<b>dcbz</b>	
<b>divw</b>	0	1	1	1	1	1																													X	<b>divw</b>



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
divw.	0	1	1	1	1	1	rD				rA				rB				0	1	1	1	1	0	1	0	1	1	1	1	X	divw.		
divwo	0	1	1	1	1	1	rD				rA				rB				1	1	1	1	1	0	1	0	1	1	0	X	divwo			
divwo.	0	1	1	1	1	1	rD				rA				rB				1	1	1	1	1	0	1	0	1	1	1	X	divwo.			
divwu	0	1	1	1	1	1	rD				rA				rB				0	1	1	1	0	0	1	0	1	1	0	X	divwu			
divwu.	0	1	1	1	1	1	rD				rA				rB				0	1	1	1	0	0	1	0	1	1	1	X	divwu.			
divwuo	0	1	1	1	1	1	rD				rA				rB				1	1	1	1	0	0	1	0	1	1	0	X	divwuo			
divwuo.	0	1	1	1	1	1	rD				rA				rB				1	1	1	1	0	0	1	0	1	1	1	X	divwuo.			
dss	dss STRM equivalent to dss STRM,0																																dss	
efdabs	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	1	0	0	1	0	0	EFX	efdabs			
efdadd	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	1	0	0	0	0	0	EFX	efdadd			
efdcfs	0	0	0	1	0	0	rD				0	0	0	0	0	rB				0	1	0	1	1	1	0	1	1	1	1	EFX	efdcfs		
efdcfsf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	0	1	1	EFX	efdcfsf			
efdcfsi	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	0	0	1	EFX	efdcfsi			
efdcfuf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	0	1	0	EFX	efdcfuf			
efdcfui	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	0	0	0	EFX	efdcfui			
efdcmpcq	0	0	0	1	0	0	crfD	/	/	rA				rB				0	1	0	1	1	1	0	1	1	1	0	EFX	efdcmpcq				
efdcmpgt	0	0	0	1	0	0	crfD	/	/	rA				rB				0	1	0	1	1	1	0	1	1	0	0	EFX	efdcmpgt				
efdcmlpt	0	0	0	1	0	0	crfD	/	/	rA				rB				0	1	0	1	1	1	0	1	1	0	1	EFX	efdcmlpt				
efdcfsf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	1	1	1	EFX	efdcfsf			
efdcfsi	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	1	0	1	EFX	efdcfsi			
efdcfsiz	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	1	0	1	0	EFX	efdcfsiz			
efdcuf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	1	1	0	EFX	efdcuf			
efdcui	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	0	1	0	0	EFX	efdcui			
efdcuiz	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	1	1	1	0	0	0	EFX	efdcuiz			
efddiv	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	1	0	1	0	0	1	EFX	efddiv			
efdmul	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	1	0	1	0	0	0	EFX	efdmul			
efdnabs	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	1	0	0	1	0	1	EFX	efdnabs			
efdneg	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	1	0	0	1	1	0	EFX	efdneg			
efdsb	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	1	0	0	0	0	1	EFX	efdsb			





Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
efdsteq	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	1	1	1	1	1	1	1	1	0	EFX	efdsteq
efdstgt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	1	1	1	1	1	1	1	0	0	EFX	efdstgt
efdstlt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	1	1	1	1	1	1	0	1	EFX	efdstlt	
efsabs	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	0	0	0	0	1	0	0	EFX	efsabs		
efsadd	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	0	0	0	0	0	0	0	EFX	efsadd		
efscfd	0	0	0	1	0	0	rD				0	0	0	0	0	rB				0	1	0	1	1	0	0	1	1	1	1	EFX	efscfd		
efscfsf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	0	1	1	EFX	efscfsf			
efscfsi	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	0	0	1	EFX	efscfsi			
efscfuf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	0	1	0	EFX	efscfuf			
efscfui	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	0	0	0	EFX	efscfui			
efscmpeq	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	1	0	EFX	efscmpeq			
efscmpgt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	0	0	EFX	efscmpgt			
efscmplt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	0	1	1	0	1	EFX	efscmplt			
efsctsf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	1	1	1	EFX	efsctsf			
efsctsi	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	1	0	1	EFX	efsctsi			
efsctsiz	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	1	0	1	0	EFX	efsctsiz			
efscruf	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	1	1	0	EFX	efscruf			
efscrtui	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	0	1	0	0	EFX	efscrtui			
efscrtuiz	0	0	0	1	0	0	rD				///				rB				0	1	0	1	1	0	1	1	0	0	0	EFX	efscrtuiz			
efsddiv	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	0	0	1	0	0	1	EFX	efsddiv			
efsmul	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	0	0	1	0	0	0	EFX	efsmul			
efsnabs	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	0	0	0	1	0	1	EFX	efsnabs			
efsneg	0	0	0	1	0	0	rD				rA				///				0	1	0	1	1	0	0	0	1	1	0	EFX	efsneg			
efssub	0	0	0	1	0	0	rD				rA				rB				0	1	0	1	1	0	0	0	0	0	1	EFX	efssub			
efststeq	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	1	1	0	EFX	efststeq			
efststgt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	1	0	0	EFX	efststgt			
efststlt	0	0	0	1	0	0	crfD		/	/	rA				rB				0	1	0	1	1	0	1	1	1	0	1	EFX	efststlt			
eqv	0	1	1	1	1	1	rD				rA				rB				0	1	0	0	0	1	1	1	0	0	0	X	eqv			
eqv.	0	1	1	1	1	1	rD				rA				rB				0	1	0	0	0	1	1	1	0	0	1	X	eqv.			



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evabs	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	EVX	evabs
evaddiw	0	1	1	1	1	1	rD			UIMM			rB			0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	EVX	evaddiw
evaddsmiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	1	0	0	1	0	0	1	0	0	1	EVX	evaddsmiaaw
evaddssiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	1	EVX	evaddssiaaw
evaddumiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	1	0	0	1	0	0	0	0	0	0	EVX	evaddumiaaw
evaddusiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evaddusiaaw
evaddw	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evaddw
evand	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1	EVX	evand
evandc	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	0	1	0	0	1	0	0	EVX	evandc	
evcmpeq	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	0	0	1	1	0	1	0	0	0	0	0	0	EVX	evcmpeq	
evcmpgts	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	0	0	1	1	0	0	0	0	1	0	0	0	1	EVX	evcmpgts
evcmpgtu	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	0	0	1	1	0	0	0	0	1	0	0	0	0	EVX	evcmpgtu
evcmplt	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	0	0	1	1	0	0	1	1	0	0	1	1	EVX	evcmplt	
evcmpltu	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	0	0	1	1	0	0	1	0	0	1	0	0	EVX	evcmpltu	
evcntlsw	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	1	1	1	1	0	0	0	EVX	evcntlsw	
evcntlzw	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	1	1	0	1	0	1	0	0	EVX	evcntlzw
evdivws	0	1	1	1	1	1	rD			rA			rB			1	0	0	1	1	0	0	0	0	1	1	0	0	0	1	1	0	EVX	evdivws
evdivwu	0	1	1	1	1	1	rD			rA			rB			1	0	0	1	1	0	0	0	0	1	1	1	1	0	0	0	EVX	evdivwu	
eveqv	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	1	0	0	1	0	0	0	1	EVX	eveqv	
evextsb	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	1	0	1	0	0	0	0	EVX	evextsb	
evextsh	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	1	0	1	1	0	0	0	EVX	evextsh	
evfsabs	0	1	1	1	1	1	rD			rA			///			0	1	0	1	0	0	0	0	0	0	1	0	0	0	0	0	EVX	evfsabs	
evfsadd	0	1	1	1	1	1	rD			rA			rB			0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evfsadd
evfscfsf	0	1	1	1	1	1	rD			///			rB			0	1	0	1	0	0	1	0	0	1	0	0	1	1	0	0	EVX	evfscfsf	
evfscfsi	0	1	1	1	1	1	rD			///			rB			0	1	0	1	0	0	1	0	0	0	0	0	0	0	1	0	EVX	evfscfsi	
evfscfuf	0	1	1	1	1	1	rD			///			rB			0	1	0	1	0	0	1	0	0	1	0	0	1	0	0	0	EVX	evfscfuf	
evfscfui	0	1	1	1	1	1	rD			///			rB			0	1	0	1	0	0	1	0	0	0	0	0	0	0	0	0	0	EVX	evfscfui
evfscmpeq	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	1	0	0	0	0	1	1	1	0	0	0	0	EVX	evfscmpeq	
evfscmpgt	0	1	1	1	1	1	crfD	/	/	/	rA			rB			0	1	0	1	0	0	0	0	1	1	0	0	0	0	0	EVX	evfscmpgt	



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evfscmplt	0	1	1	1	1	1	crfD		/	/	rA						rB					0	1	0	1	0	0	0	1	1	0	1	EVX	evfscmplt
evfsctsf	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	0	1	1	1	EVX	evfsctsf		
evfsctsi	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	0	1	0	1	EVX	evfsctsi		
evfsctsiz	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	1	0	1	0	EVX	evfsctsiz		
evfsctuf	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	0	1	1	0	EVX	evfsctuf		
evfsctui	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	0	1	0	0	EVX	evfsctui		
evfsctuiiz	0	1	1	1	1	1	rD		///						rB					0	1	0	1	0	0	1	1	0	0	0	EVX	evfsctuiiz		
evfsdiv	0	1	1	1	1	1	rD		rA						rB					0	1	0	1	0	0	0	1	0	0	1	EVX	evfsdiv		
evfsmul	0	1	1	1	1	1	rD		rA						rB					0	1	0	1	0	0	0	1	0	0	0	EVX	evfsmul		
evfsnabs	0	1	1	1	1	1	rD		rA						///					0	1	0	1	0	0	0	0	1	0	1	EVX	evfsnabs		
evfsneg	0	1	1	1	1	1	rD		rA						///					0	1	0	1	0	0	0	0	1	1	0	EVX	evfsneg		
evfssub	0	1	1	1	1	1	rD		rA						rB					0	1	0	1	0	0	0	0	0	0	1	EVX	evfssub		
evfststeq	0	1	1	1	1	1	crfD		/	/	rA						rB					0	1	0	1	0	0	1	1	1	1	0	EVX	evfststeq
evfststgt	0	1	1	1	1	1	crfD		/	/	rA						rB					0	1	0	1	0	0	1	1	1	0	0	EVX	evfststgt
evfststlt	0	1	1	1	1	1	crfD		/	/	rA						rB					0	1	0	1	0	0	1	1	1	0	1	EVX	evfststlt
evlidd	0	1	1	1	1	1	rD		rA						UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	0	0	1	EVX	evlidd		
evliddx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	0	0	0	0	EVX	evliddx		
evldh	0	1	1	1	1	1	rD		rA						UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	1	0	1	EVX	evldh		
evldhx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	0	1	0	0	EVX	evldhx		
evldw	0	1	1	1	1	1	rD		rA						UIMM <sup>(8)</sup>					0	1	1	0	0	0	0	0	0	1	1	EVX	evldw		
evldwx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	0	0	1	0	EVX	evldwx		
evlhhesplat	0	1	1	1	1	1	rD		rA						UIMM <sup>(9)</sup>					0	1	1	0	0	0	0	1	0	0	1	EVX	evlhhesplat		
evlhhesplatx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	1	0	0	0	EVX	evlhhesplatx		
evlhhosspat	0	1	1	1	1	1	rD		rA						UIMM <sup>(9)</sup>					0	1	1	0	0	0	0	1	1	1	1	EVX	evlhhosspat		
evlhhosspatx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	1	1	1	0	EVX	evlhhosspatx		
evlhhuspat	0	1	1	1	1	1	rD		rA						UIMM <sup>(9)</sup>					0	1	1	0	0	0	0	1	1	0	1	EVX	evlhhuspat		
evlhhuspatx	0	1	1	1	1	1	rD		rA						rB					0	1	1	0	0	0	0	1	1	0	0	EVX	evlhhuspatx		
evlwhe	0	1	1	1	1	1	rD		rA						UIMM <sup>(10)</sup>					0	1	1	0	0	0	1	0	0	0	1	EVX	evlwhe		



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evlwhex	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	0	0	0	1	0	0	0	0	EVX	evlwhex	
evlwhos	0	1	1	1	1	1	rD			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	0	1	0	1	1	1	1	1	1	1	1	EVX	evlwhos	
evlwhosx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	1	1	1	0	1	1	1	0	EVX	evlwhosx	
evlwhou	0	1	1	1	1	1	rD			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	0	1	0	1	0	1	0	1	0	1	1	EVX	evlwhou	
evlwhoux	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	1	0	1	0	0	1	0	0	EVX	evlwhoux	
evlwhsplat	0	1	1	1	1	1	rD			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	0	1	1	1	1	0	1	1	1	0	1	EVX	evlwhsplat	
evlwhsplatx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	1	1	1	0	0	1	1	1	0	0	EVX	evlwhsplatx
evlwwsplat	0	1	1	1	1	1	rD			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	0	1	1	1	0	0	1	1	0	0	1	EVX	evlwwsplat	
evlwwsplatx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	1	1	0	0	0	1	1	0	0	0	EVX	evlwwsplatx
evmergehi	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	1	1	1	0	0	1	1	0	0	EVX	evmergehi	
evmergehilo	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	1	1	1	1	0	1	1	1	0	EVX	evmergehilo	
evmergeelo	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	1	1	1	0	1	1	0	1	1	EVX	evmergeelo	
evmergeलोhi	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	1	1	1	1	1	1	1	1	1	EVX	evmergeलोhi	
evmhegsmfaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	0	1	0	1	1	1	1	1	EVX	evmhegsmfaa	
evmhegsmfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	1	1	1	1	1	EVX	evmhegsmfan	
evmhegsmiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	0	0	0	1	0	0	1	1	EVX	evmhegsmiaa	
evmhegsmian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	0	1	1	1	1	EVX	evmhegsmian	
evmhegumiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	0	0	0	1	0	0	0	0	EVX	evmhegumiaa	
evmhegumian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	0	1	0	0	0	0	0	0	EVX	evmhegumian	
evmhesmf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	0	1	0	1	1	1	1	EVX	evmhesmf	
evmhesmfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	0	1	0	1	1	1	1	1	EVX	evmhesmfa	
evmhesmfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	0	1	0	1	1	1	1	EVX	evmhesmfaaw	
evmhesmfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	0	1	0	1	1	1	1	1	EVX	evmhesmfanw	
evmhesmi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	0	0	0	1	0	0	1	EVX	evmhesmi	
evmhesmia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	0	0	0	1	0	0	1	1	EVX	evmhesmia	
evmhesmiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	0	0	0	1	0	0	1	EVX	evmhesmiaaw	
evmhesmianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	0	0	0	1	0	0	1	1	EVX	evmhesmianw	
evmhessf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	EVX	evmhessf	



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhessfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	0	0	1	1	1	1	1	1	EVX	evmhessfa		
evmhessfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	1	1	1	EVX	evmhessfaaw		
evmhessfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	1	1	1	EVX	evmhessfanw			
evmhessiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	1	1	1	EVX	evmhessiaaw		
evmhessianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	1	1	1	EVX	evmhessianw		
evmheumi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	EVX	evmheumi		
evmheumia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	0	0	0	0	0	0	0	EVX	evmheumia		
evmheumiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	0	0	0	0	0	0	EVX	evmheumiaaw		
evmheumianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	0	0	0	0	0	0	0	EVX	evmheumianw		
evmheusiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evmheusiaaw		
evmheusianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evmheusianw		
evmhogsmfaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	1	1	1	1	1	EVX	evmhogsmfaa			
evmhogsmfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	1	1	1	1	1	EVX	evmhogsmfan			
evmhogsmiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	0	1	1	1	1	EVX	evmhogsmiaa			
evmhogsmian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	0	1	1	1	1	EVX	evmhogsmian			
evmhogumiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	0	0	0	0	0	0	EVX	evmhogumiaa		
evmhogumian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	1	0	0	0	0	0	EVX	evmhogumian		
evmhosmf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	1	1	1	1	1	EVX	evmhosmf			
evmhosmfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	1	1	1	1	1	1	EVX	evmhosmfa			
evmhosmfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	1	1	1	1	1	EVX	evmhosmfaaw			
evmhosmfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	1	1	1	1	1	1	EVX	evmhosmfanw			
evmhosmi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	1	0	1	1	1	EVX	evmhosmi			
evmhosmia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	1	0	1	1	1	1	EVX	evmhosmia			
evmhosmiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	1	1	0	1	1	1	1	EVX	evmhosmiaaw			
evmhosmianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	1	0	1	1	1	1	EVX	evmhosmianw			
evmhossf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	0	1	1	1	1	1	EVX	evmhossf			
evmhossfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	0	1	1	1	1	1	1	EVX	evmhossfa			
evmhossfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	1	1	1	1	1	EVX	evmhossfaaw			
evmhossfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	1	1	1	1	1	EVX	evmhossfanw			



**Table 271. Instructions sorted by mnemonic (binary) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhossiaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	1	0	1	EVX	evmhossiaaw		
evmhossianw	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	1	0	1	EVX	evmhossianw		
evmhoumi	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	0	0	0	0	1	1	0	0	EVX	evmhoumi			
evmhoumia	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	0	1	1	1	0	0	EVX	evmhoumia				
evmhoumiaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	0	0	0	0	1	1	0	0	EVX	evmhoumiaaw				
evmhoumianw	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	0	0	0	0	1	1	0	0	EVX	evmhoumianw				
evmhousiaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	0	0	0	0	0	0	1	0	0	EVX	evmhousiaaw			
evmhousianw	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	0	0	0	0	0	0	1	0	0	EVX	evmhousianw			
evmr	evmr rD,rA						equivalent to						evor rD,rA,rA																				evmr	
evmra	0	1	1	1	1	1		rD			rA			///			1	0	0	1	1	0	0	0	0	1	0	0	EVX	evmra				
evmwhgsmfaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	1	0	1	1	1	1	1	1	EVX	evmwhgsmfaa			
evmwhgsmfan	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	1	1	1	1	1	EVX	evmwhgsmfan			
evmwhgsmiaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	1	0	1	1	0	1	EVX	evmwhgsmiaa					
evmwhgsmian	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	EVX	evmwhgsmian					
evmwhgssfaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	1	0	0	1	1	1	1	EVX	evmwhgssfaa				
evmwhgssfan	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	0	1	1	1	1	EVX	evmwhgssfan				
evmwhgumiaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	1	0	1	1	0	0	EVX	evmwhgumiaa					
evmwhgumian	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	1	0	0	EVX	evmwhgumian					
evmwhsmf	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	0	0	1	1	1	1	1	EVX	evmwhsmf				
evmwhsmfa	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	1	0	1	1	1	1	1	EVX	evmwhsmfa				
evmwhsmfaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	0	1	1	1	1	1	EVX	evmwhsmfaaw				
evmwhsmfanw	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	0	1	1	1	1	1	EVX	evmwhsmfanw				
evmwhsmi	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	0	0	1	1	0	1	EVX	evmwhsmi					
evmwhsmia	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	1	0	1	1	0	1	EVX	evmwhsmia					
evmwhsmiaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	0	1	1	0	1	EVX	evmwhsmiaaw					
evmwhsmianw	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	0	1	1	0	1	EVX	evmwhsmianw					
evmwhssf	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	0	0	0	1	1	1	1	EVX	evmwhssf				
evmwhssfa	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	1	1	0	0	1	1	1	1	EVX	evmwhssfa				
evmwhssfaaw	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	0	0	1	1	1	1	EVX	evmwhssfaaw				





**Table 271. Instructions sorted by mnemonic (binary) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evwsmfan	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	1	0	1	1	EVX	evwsmfan	
evwsmi	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	0	1	1	1	0	0	0	1	EVX	evwsmi		
evwsmia	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	1	1	1	0	0	0	1	EVX	evwsmia			
evwsmiaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	1	1	1	0	0	0	1	EVX	evwsmiaa			
evwsmian	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	1	0	0	0	1	EVX	evwsmian			
evwssf	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	0	1	0	0	0	1	1	EVX	evwssf			
evwssfa	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	1	1	0	0	1	1	EVX	evwssfa				
evwssfaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	1	0	0	0	1	1	EVX	evwssfaa				
evwssfan	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	0	0	1	1	EVX	evwssfan					
evwumi	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	0	1	1	0	0	0	EVX	evwumi				
evwumia	0	1	1	1	1	1		rD			rA			rB			1	0	0	0	0	1	1	1	1	0	0	0	EVX	evwumia				
evwumiaa	0	1	1	1	1	1		rD			rA			rB			1	0	1	0	1	0	1	1	1	0	0	0	EVX	evwumiaa				
evwumian	0	1	1	1	1	1		rD			rA			rB			1	0	1	1	1	0	1	1	0	0	0	EVX	evwumian					
evnand	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	0	1	1	1	1	0	EVX	evnand				
evneg	0	1	1	1	1	1		rD			rA			///			0	1	0	0	0	0	0	0	1	0	0	1	EVX	evneg				
evnor	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	0	1	1	0	0	0	EVX	evnor				
evnot	evnot rD,rA						equivalent to						evnor rD,rA,rA												evnot									
evor	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	0	1	0	1	1	1	EVX	evor				
evorc	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	0	1	1	0	1	1	EVX	evorc				
evrlw	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	1	0	1	0	0	0	EVX	evrlw				
evrlwi	0	1	1	1	1	1		rD			rA			UIMM			0	1	0	0	0	0	1	0	1	0	1	0	EVX	evrlwi				
evrndw	0	1	1	1	1	1		rD			rA			UIMM			0	1	0	0	0	0	0	0	1	1	0	0	EVX	evrndw				
evsel	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	1	1	1	crfS			EVX	evsel					
evslw	0	1	1	1	1	1		rD			rA			rB			0	1	0	0	0	0	1	0	0	1	0	0	EVX	evslw				
evslwi	0	1	1	1	1	1		rD			rA			UIMM			0	1	0	0	0	0	1	0	0	1	1	0	EVX	evslwi				
evsplatfi	0	1	1	1	1	1		rD			SIMM			///			0	1	0	0	0	0	1	0	1	0	1	1	EVX	evsplatfi				
evsplatli	0	1	1	1	1	1		rD			SIMM			///			0	1	0	0	0	0	1	0	1	0	0	1	EVX	evsplatli				
evsrwis	0	1	1	1	1	1		rD			rA			UIMM			0	1	0	0	0	0	1	0	0	0	1	1	EVX	evsrwis				
evsrwiu	0	1	1	1	1	1		rD			rA			UIMM			0	1	0	0	0	0	1	0	0	0	1	0	EVX	evsrwiu				





Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evsrws	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	0	0	0	1	0	0	0	0	1	EVX	evsrws	
evsrwu	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	EVX	evsrwu	
evstdd	0	1	1	1	1	1	rD			rA			UIMM <sup>(8)</sup>			0	1	1	0	0	1	0	0	0	0	1	0	0	0	0	1	EVX	evstdd	
evstddx	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	EVX	evstddx	
evstdh	0	1	1	1	1	1	rS			rA			UIMM <sup>(8)</sup>			0	1	1	0	0	1	0	0	1	0	0	1	0	1	0	1	EVX	evstdh	
evstdhx	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	0	0	1	0	0	1	0	0	0	0	EVX	evstdhx	
evstdw	0	1	1	1	1	1	rS			rA			UIMM <sup>(8)</sup>			0	1	1	0	0	1	0	0	0	1	0	0	0	1	1	EVX	evstdw		
evstdwx	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	0	0	0	1	0	0	0	1	0	EVX	evstdwx		
evstwhe	0	1	1	1	1	1	rS			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	1	1	0	0	1	1	0	0	0	1	EVX	evstwhe		
evstwhex	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	1	0	0	0	0	0	0	0	0	EVX	evstwhex		
evstwho	0	1	1	1	1	1	rS			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	1	1	0	1	0	1	0	1	0	1	EVX	evstwho		
evstwhox	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	1	0	1	0	1	0	0	0	0	EVX	evstwhox		
evstwwe	0	1	1	1	1	1	rS			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	1	1	1	1	0	0	1	0	0	1	EVX	evstwwe		
evstwwex	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	1	1	1	0	0	0	0	0	0	EVX	evstwwex		
evstwwo	0	1	1	1	1	1	rS			rA			UIMM <sup>(10)</sup>			0	1	1	0	0	1	1	1	1	1	0	1	0	1	EVX	evstwwo			
evstwwox	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	1	1	1	1	1	0	0	0	0	EVX	evstwwox		
evsubfsmiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	1	0	1	0	1	1	1	1	EVX	evsubfsmiaaw		
evsubfssiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	1	1	1	1	1	EVX	evsubfssiaaw		
evsubfumiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	1	0	1	0	1	0	0	0	EVX	evsubfumiaaw		
evsubfusiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	1	0	0	0	0	EVX	evsubfusiaaw		
evsubfw	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	0	0	1	0	0	0	0	0	EVX	evsubfw		
evsubifw	0	1	1	1	1	1	rD			UIMM			rB			0	1	0	0	0	0	0	0	0	0	1	1	0	0	0	EVX	evsubifw		
evsubiw	evsubiw rD,rB,UIMM equivalent to evsubifw rD,UIMM,rB																															evsubiw		
evsubw	evsubw rD,rB,rA equivalent to evsubfw rD,rA,rB																															evsubw		
evxor	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	1	0	1	1	1	0	0	0	0	EVX	evxor		
extlwi	extlwi rA,rS,n,b (n > 0) equivalent to rlwinm rA,rS,b,0,n-1																															extlwi		
extrwi	extrwi rA,rS,n,b (n > 0) equivalent to rlwinm rA,rS,b+n,32-n,31																															extrwi		
extsb	0	1	1	1	1	1	rS			rA			///			1	1	1	0	1	1	1	0	1	0	0	0	0	0	0	X	extsb		



**Table 271. Instructions sorted by mnemonic (binary) (continued)**

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
extsb.	0	1	1	1	1	1	rS			rA			///			1	1	1	0	1	1	1	0	1	0	1					X	extsb.		
extsh	0	1	1	1	1	1	rS			rA			///			1	1	1	0	0	1	1	0	1	0	0					X	extsh		
extsh.	0	1	1	1	1	1	rS			rA			///			1	1	1	0	0	1	1	0	1	0	1					X	extsh.		
fres <sup>(6)</sup>	1	1	1	0	1	1	frD			///			frB			///			1	1	0	0	0	0	0					A	fres			
fres. <sup>(6)</sup>	1	1	1	0	1	1	frD			///			frB			///			1	1	0	0	0	0	1					A	fres.			
fsel <sup>(6)</sup>	1	1	1	1	1	1	frD			frA			frB			frC			1	0	1	1	1	0					A	fsel				
fsel. <sup>(6)</sup>	1	1	1	1	1	1	frD			frA			frB			frC			1	0	1	1	1	1					A	fsel.				
icbi	0	1	1	1	1	1	///			rA			rB			1	1	1	1	0	1	0	1	1	0	/				X	icbi			
icblc	0	1	1	1	1	1	CT			rA			rB			0	0	1	1	1	0	0	1	1	0	0				X	icblc			
icbt	0	1	1	1	1	1	CT			rA			rB			0	0	0	0	0	1	0	1	1	0	/				X	icbt			
icbtlc	0	1	1	1	1	1	CT			rA			rB			0	1	1	1	1	0	0	1	1	0	0				X	icbtlc			
inslwi	inslwi rA,rS,n,b (n > 0)						equivalent to						rlwimi rA,rS,32 - b,b,(b + n) - 1												inslwi									
insrwi	insrwi rA,rS,n,b (n > 0)						equivalent to						rlwimi rA,rS,32 - (b + n),b,(b + n) - 1												insrwi									
isel	0	1	1	1	1	1	rD			rA			rB			crb			0	1	1	1	1	0					X	isel				
iseleq	iseleq rD,rA,rB						equivalent to						isel rD,rA,rB,2												iseleq									
iselgt	iselgt rD,rA,rB						equivalent to						isel rD,rA,rB,1												iselgt									
isellt	isellt rD,rA,rB						equivalent to						isel rD,rA,rB,0												isellt									
isync	0	1	0	0	1	1	///			///			0	0	1	0	0	1	0	1	1	0	1	1	0	/				XL	isync			
la	la rD,d(rA)						equivalent to						addi rD,rA,d												la									
lbz	1	0	0	0	1	0	rD			rA			D						D	lbz														
lbzu	1	0	0	0	1	1	rD			rA			D						D	lbzu														
lbzux	0	1	1	1	1	1	rD			rA			rB			0	0	0	1	1	1	0	1	1	1	/				X	lbzux			
lbzx	0	1	1	1	1	1	rD			rA			rB			0	0	0	1	0	1	0	1	1	1	/				X	lbzx			
lha	1	0	1	0	1	0	rD			rA			D						D	lha														
lhau	1	0	1	0	1	1	rD			rA			D						D	lhau														
lhaux	0	1	1	1	1	1	rD			rA			rB			0	1	0	1	1	1	0	1	1	1	/				X	lhaux			
lhax	0	1	1	1	1	1	rD			rA			rB			0	1	0	1	0	1	0	1	1	1	/				X	lhax			
lhbrx	0	1	1	1	1	1	rD			rA			rB			1	1	0	0	0	1	0	1	1	0	/				X	lhbrx			



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
lhz	1	0	1	0	0	0	rD				rA				D												D	lhz						
lhzu	1	0	1	0	0	1	rD				rA				D												D	lhzu						
lhzux	0	1	1	1	1	1	rD				rA				rB		0	1	0	0	1	1	0	1	1	1	1	/	X	lhzux				
lhzx	0	1	1	1	1	1	rD				rA				rB		0	1	0	0	0	1	0	1	1	1	/	X	lhzx					
li	li rD,value						equivalent to						addi rD,0,value																		li			
lis	lis rD,value						equivalent to						addis rD,0,value																		lis			
lmw	1	0	1	1	1	0	rD				rA				D												D	lmw						
lwarx	0	1	1	1	1	1	rD				rA				rB		0	0	0	0	0	1	0	1	0	0	/	X	lwarx					
lwbrx	0	1	1	1	1	1	rD				rA				rB		1	0	0	0	0	1	0	1	1	0	/	X	lwbrx					
lwz	1	0	0	0	0	0	rD				rA				D												D	lwz						
lwzu	1	0	0	0	0	1	rD				rA				D												D	lwzu						
lwzux	0	1	1	1	1	1	rD				rA				rB		0	0	0	0	1	1	0	1	1	1	/	X	lwzux					
lwzx	0	1	1	1	1	1	rD				rA				rB		0	0	0	0	0	1	0	1	1	1	/	X	lwzx					
mbar	0	1	1	1	1	1	MO				///				1	1	0	1	0	1	0	1	1	0	/	X	mbar							
mcrf	0	1	0	0	1	1	crfD		//	crfS		///				0	0	0	0	0	0	0	0	0	0	0	/	XL	mcrf					
mcrxr	0	1	1	1	1	1	crfD		///				1	0	0	0	0	0	0	0	0	0	0	0	0	/	X	mcrxr						
mfcrr	mfcrr rS						equivalent to						mfcrr 0xFF,rS																		mfcrr			
mfcrr	0	1	1	1	1	1	rD				///				0	0	0	0	0	1	0	0	1	1	/	X	mfcrr							
mfdcrr	0	1	1	1	1	1	rD				DCRN5-9				DCRN0-4				0	1	0	1	0	0	0	1	1	/	XF	mfdcrr				
mfmsrr <sup>(7)</sup>	0	1	1	1	1	1	rD				///				0	0	0	1	0	1	0	0	1	1	/	X	mfmsrr							
mfpmrr	0	1	1	1	1	1	rD				PMRN5-9				PMRN0-4				0	1	0	1	0	0	1	1	1	0	0	XF	mfpmrr			
mfregname	mfregname rD						equivalent to						mfspr rD,SPRn																		mfregname			
mfspr <sup>(11)</sup>	0	1	1	1	1	1	rD				SPR[5-9]				SPR[0-4]				0	1	0	1	0	1	0	0	1	1	/	XF	mfspr			
mr	mr rA,rS						equivalent to						or rA,rS,rS																		mr			
msync	0	1	1	1	1	1	///				1	0	0	1	0	1	0	1	0	1	1	0	/	X	msync									
mtcrr	mtcrr rS						equivalent to						mtcrr 0xFF,rS																		mtcrr			
mtcrrf	0	1	1	1	1	1	rS				/	CRM				/	0	0	1	0	0	1	0	0	0	0	/	XF	mtcrrf					
mtdcrr	0	1	1	1	1	1	rS				DCRN5-9				DCRN0-4				0	1	1	1	0	0	0	1	1	/	XF	mtdcrr				
mtmsrr <sup>(7)</sup>	0	1	1	1	1	1	rS				///				0	0	1	0	0	1	0	0	1	0	/	X	mtmsrr							



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic														
mtpmr	0	1	1	1	1	1	rS				PMRN5-9				PMRN0-4				0	1	1	1	0	0	1	1	1	0	0	XFX	mtpmr																	
mtregname	mtregname rS						equivalent to																														mtspr SPRn rS						mtregname					
mtspr <sup>(11)</sup>	0	1	1	1	1	1	rS				SPR[5-9]				SPR[0-4]				0	1	1	1	0	1	0	0	1	1	/	XFX	mtspr																	
mulhw	0	1	1	1	1	1	rD				rA				rB				/	0	0	1	0	0	1	0	1	1	0	X	mulhw																	
mulhw.	0	1	1	1	1	1	rD				rA				rB				/	0	0	1	0	0	1	0	1	1	1	X	mulhw.																	
mulhwu	0	1	1	1	1	1	rD				rA				rB				/	0	0	0	0	0	1	0	1	1	0	X	mulhwu																	
mulhwu.	0	1	1	1	1	1	rD				rA				rB				/	0	0	0	0	0	1	0	1	1	1	X	mulhwu.																	
mulli	0	0	0	1	1	1	rD				rA				SIMM										D	mulli																						
mullw	0	1	1	1	1	1	rD				rA				rB				0	0	1	1	1	0	1	0	1	1	0	X	mullw																	
mullw.	0	1	1	1	1	1	rD				rA				rB				0	0	1	1	1	0	1	0	1	1	1	X	mullw.																	
mullwo	0	1	1	1	1	1	rD				rA				rB				1	0	1	1	1	0	1	0	1	1	0	X	mullwo																	
mullwo.	0	1	1	1	1	1	rD				rA				rB				1	0	1	1	1	0	1	0	1	1	1	X	mullwo.																	
nand	0	1	1	1	1	1	rS				rA				rB				0	1	1	1	0	1	1	1	0	0	0	X	nand																	
nand.	0	1	1	1	1	1	rS				rA				rB				0	1	1	1	0	1	1	1	0	0	1	X	nand.																	
neg	0	1	1	1	1	1	rD				rA				///				0	0	0	1	1	0	1	0	0	0	0	X	neg																	
neg.	0	1	1	1	1	1	rD				rA				///				0	0	0	1	1	0	1	0	0	0	1	X	neg.																	
nego	0	1	1	1	1	1	rD				rA				///				1	0	0	1	1	0	1	0	0	0	0	X	nego																	
nego.	0	1	1	1	1	1	rD				rA				///				1	0	0	1	1	0	1	0	0	0	1	X	nego.																	
nop	nop						equivalent to																														ori 0,0,0						nop					
nor	0	1	1	1	1	1	rS				rA				rB				0	0	0	1	1	1	1	1	1	0	0	0	X	nor																
nor.	0	1	1	1	1	1	rS				rA				rB				0	0	0	1	1	1	1	1	1	0	0	1	X	nor.																
not	not rA,rS						equivalent to																														nor rA,rS,rS						not					
or	0	1	1	1	1	1	rS				rA				rB				0	1	1	0	1	1	1	1	1	0	0	0	X	or																
or.	0	1	1	1	1	1	rS				rA				rB				0	1	1	0	1	1	1	1	1	0	0	1	X	or.																
orc	0	1	1	1	1	1	rS				rA				rB				0	1	1	0	0	1	1	1	0	0	0	X	orc																	
orc.	0	1	1	1	1	1	rS				rA				rB				0	1	1	0	0	1	1	1	0	0	1	X	orc.																	
ori	0	1	1	0	0	0	rS				rA				UIMM										D	ori																						
oris	0	1	1	0	0	1	rS				rA				UIMM										D	oris																						
rfci	0	1	0	0	1	1	///				///				///				0	0	0	0	1	1	0	0	1	1	/	XL	rfci																	



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
rfdi <sup>(7)</sup>	0	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	1	1	0	X	rfdi	
rfi <sup>(7)</sup>	0	1	0	0	1	1	///															0	0	0	0	1	1	0	0	1	0	/	XL	rfi
rfmci <sup>(7)</sup>	0	1	0	0	1	1	///															0	0	0	0	1	0	0	1	1	0	/	XL	rfmci
rlwimi	0	1	0	1	0	0	rS			rA			SH			MB			ME			0	M	rlwimi										
rlwimi.	0	1	0	1	0	0	rS			rA			SH			MB			ME			1	M	rlwimi.										
rlwinm	0	1	0	1	0	1	rS			rA			SH			MB			ME			0	M	rlwinm										
rlwinm.	0	1	0	1	0	1	rS			rA			SH			MB			ME			1	M	rlwinm.										
rlwnm	0	1	0	1	1	1	rS			rA			rB			MB			ME			0	M	rlwnm										
rlwnm.	0	1	0	1	1	1	rS			rA			rB			MB			ME			1	M	rlwnm.										
rotlw	rotlw rA,rS,rB						equivalent to															rlwnm rA,rS,rB,0,31								rotlw				
rotlwi	rotlwi rA,rS,n						equivalent to															rlwinm rA,rS,n,0,31								rotlwi				
rotrwi	rotrwi rA,rS,n						equivalent to															rlwinm rA,rS,32 - n,0,31								rotrwi				
sc	0	1	0	0	0	1	///															1	/	SC	sc									
slw	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	1	1	0	0	0	0	X	slw						
slw.	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	1	1	0	0	0	1	X	slw.						
slwi	slwi rA,rS,n (n < 32)						equivalent to															rlwinm rA,rS,n,0,31 - n								slwi				
sraw	0	1	1	1	1	1	rS			rA			rB			1	1	0	0	0	1	1	0	0	0	0	0	X	sraw					
sraw.	0	1	1	1	1	1	rS			rA			rB			1	1	0	0	0	1	1	0	0	0	1	X	sraw.						
srawi	0	1	1	1	1	1	rS			rA			SH			1	1	0	0	1	1	1	0	0	0	0	X	srawi						
srawi.	0	1	1	1	1	1	rS			rA			SH			1	1	0	0	1	1	1	0	0	0	1	X	srawi.						
srw	0	1	1	1	1	1	rS			rA			rB			1	0	0	0	0	1	1	0	0	0	0	X	srw						
srw.	0	1	1	1	1	1	rS			rA			rB			1	0	0	0	0	1	1	0	0	0	1	X	srw.						
srwi	srwi rA,rS,n (n < 32)						equivalent to															rlwinm rA,rS,32 - n,n,31								srwi				
stb	1	0	0	1	1	0	rS			rA			D															D	stb					
stbu	1	0	0	1	1	1	rS			rA			D															D	stbu					
stbux	0	1	1	1	1	1	rS			rA			rB			0	0	1	1	1	1	0	1	1	1	0	X	stbux						
stbx	0	1	1	1	1	1	rS			rA			rB			0	0	1	1	0	1	0	1	1	1	0	X	stbx						
sth	1	0	1	1	0	0	rS			rA			D															D	sth					
sthbrx	0	1	1	1	1	1	rS			rA			rB			1	1	1	0	0	1	0	1	1	0	/	X	sthbrx						



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic	
sth	1	0	1	1	0	1	rS			rA			D															D	sth						
sthux	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	1	1	0	1	1	0	1	1	1	1	/	X	sthux			
sthx	0	1	1	1	1	1	rS			rA			rB			0	1	1	0	0	1	0	1	1	1	1	/	X	sthx						
stmw	1	0	1	1	1	1	rS			rA			D															D	stmw						
stw	1	0	0	1	0	0	rS			rA			D															D	stw						
stwbrx	0	1	1	1	1	1	rS			rA			rB			1	0	1	0	0	1	0	1	1	1	0	/	X	stwbrx						
stwcx.	0	1	1	1	1	1	rS			rA			rB			0	0	1	0	0	1	0	1	1	1	0	1	X	stwcx.						
stwu	1	0	0	1	0	1	rS			rA			D															D	stwu						
stwux	0	1	1	1	1	1	rS			rA			rB			0	0	1	0	1	1	0	1	1	1	/	D	stwux							
stwx	0	1	1	1	1	1	rS			rA			rB			0	0	1	0	0	1	0	1	1	1	/	D	stwx							
sub	sub rD,rA,rB						equivalent to						subf rD,rB,rA																						sub
subc	subc rD,rA,rB						equivalent to						subfc rD,rB,rA																						subc
subf	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	1	0	1	0	0	0	0	0	X	subf						
subf.	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	1	0	1	0	0	0	1	X	subf.							
subfc	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	1	0	0	0	0	X	subfc							
subfc.	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	1	0	0	0	1	X	subfc.							
subfco	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	1	0	0	0	0	X	subfco							
subfco.	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	1	0	0	0	1	X	subfco.							
subfe	0	1	1	1	1	1	rD			rA			rB			0	0	1	0	0	0	1	0	0	0	0	X	subfe							
subfe.	0	1	1	1	1	1	rD			rA			rB			0	0	1	0	0	0	1	0	0	0	1	X	subfe.							
subfeo	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	1	0	0	0	0	X	subfeo							
subfeo.	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	1	0	0	0	1	X	subfeo.							
subfic	0	0	1	0	0	0	rD			rA			SIMM															D	subfic						
subfme	0	1	1	1	1	1	rD			rA			///			0	0	1	1	1	0	1	0	0	0	0	X	subfme							
subfme.	0	1	1	1	1	1	rD			rA			///			0	0	1	1	1	0	1	0	0	0	1	X	subfme.							
subfmeo	0	1	1	1	1	1	rD			rA			///			1	0	1	1	1	0	1	0	0	0	0	X	subfmeo							
subfmeo.	0	1	1	1	1	1	rD			rA			///			1	0	1	1	1	0	1	0	0	0	1	X	subfmeo.							
subfo	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	0	0	0	0	X	subfo							
subfo.	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	0	0	0	1	X	subfo.							



Table 271. Instructions sorted by mnemonic (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
subfze	0	1	1	1	1	1	rD				rA				///				0	0	1	1	0	0	1	0	0	0	0	0	0	X	subfze	
subfze.	0	1	1	1	1	1	rD				rA				///				0	0	1	1	0	0	1	0	0	0	0	1	X	subfze.		
subfzeo	0	1	1	1	1	1	rD				rA				///				1	0	1	1	0	0	1	0	0	0	0	X	subfzeo			
subfzeo.	0	1	1	1	1	1	rD				rA				///				1	0	1	1	0	0	1	0	0	0	1	X	subfzeo.			
subi	subi rD,rA,value equivalent to										addi rD,rA,-value																		subi					
subic	subic rD,rA,value equivalent to										addic rD,rA,-value																		subic					
subic.	subic. rD,rA,value equivalent to										addic. rD,rA,-value																		subic.					
subis	subis rD,rA,value equivalent to										addis rD,rA,-value																		subis					
tlbie <sup>(6),(7)</sup>	0	1	1	1	1	1	///				///				rB				0	1	0	0	1	1	0	0	1	0	0	X	tlbie			
tlbivax	0	1	1	1	1	1	///				rA				rB				1	1	0	0	0	1	0	0	1	0	/	X	tlbivax			
tlbre	0	1	1	1	1	1	/// <sup>(12)</sup>										1	1	1	0	1	1	0	0	1	0	/	X	tlbre					
tlbsx	0	1	1	1	1	1	/// <sup>(12)</sup>				rA				rB				1	1	1	0	0	1	0	0	1	0	/12	X	tlbsx			
tlbsync <sup>(6),(7)</sup>	0	1	1	1	1	1	///				///				///				1	0	0	0	1	1	0	1	1	0	/	X	tlbsync			
tlbwe	0	1	1	1	1	1	/// <sup>(12)</sup>										1	1	1	1	0	1	0	0	1	0	/	X	tlbwe					
tw	0	1	1	1	1	1	TO				rA				rB				0	0	0	0	0	0	0	1	0	0	/	X	tw			
tweq	tweq rA,SIMM equivalent to										tw 4,rA,SIMM																		tweq					
tweqi	tweqi rA,SIMM equivalent to										twi 4,rA,SIMM																		tweqi					
twge	twge rA,SIMM equivalent to										tw 12,rA,SIMM																		twge					
twgei	twgei rA,SIMM equivalent to										twi 12,rA,SIMM																		twgei					
twgt	twgt rA,SIMM equivalent to										tw 8,rA,SIMM																		twgt					
twgti	twgti rA,SIMM equivalent to										twi 8,rA,SIMM																		twgti					
twi	0	0	0	0	1	1	TO				rA				SIMM										D	twi								
twle	twle rA,SIMM equivalent to										tw 20,rA,SIMM																		twle					
twlei	twlei rA,SIMM equivalent to										twi 20,rA,SIMM																		twlei					
twlge	twlge rA,SIMM equivalent to										tw 12,rA,SIMM																		twlge					
twlgei	twlgei rA,SIMM equivalent to										twi 12,rA,SIMM																		twlgei					
twlgt	twlgt rA,SIMM equivalent to										tw 1,rA,SIMM																		twlgt					
twlgti	twlgti rA,SIMM equivalent to										twi 1,rA,SIMM																		twlgti					





7. Supervisor-level instruction.
8.  $d = \text{UIMM} * 8$
9. Access level is determined by whether the SPR is defined as a user or supervisor level SPR.
10.  $d = \text{UIMM} * 4$
11. Access level is determined by whether the SPR is defined as a user- or supervisor-level SPR.
12. This field is defined as allocated by the Book E architecture, for possible use in an implementation.



## A.4 Instructions sorted by opcode (binary)

[Table 272](#) lists instructions by opcode, shown in binary.



Table 272. Instructions sorted by opcode (binary)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
rfdi <sup>(1)</sup>	0	1	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	1	1	0	X	rfdi
twi	0	0	0	0	1	1	TO			rA			SIMM															D	twi					
brinc	0	0	0	1	0	0	rD			rA			rB			0	1	0	0	0	0	0	0	0	1	1	1	1	1	1	1	EVX	brinc	
efdabs	0	0	0	1	0	0	rD			rA			///			0	1	0	1	1	1	1	0	0	1	0	0	1	0	0	0	EFX	efdabs	
efdadd	0	0	0	1	0	0	rD			rA			rB			0	1	0	1	1	1	1	0	0	0	0	0	0	0	0	0	EFX	efdadd	
efdcfs	0	0	0	1	0	0	rD			0	0	0	0	0	0	rB			0	1	0	1	1	1	1	0	1	1	1	1	1	1	EFX	efdcfs
efdcfsf	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	0	1	1	1	1	1	1	EFX	efdcfsf	
efdcfsi	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	0	0	0	1	1	1	1	EFX	efdcfsi	
efdcfuf	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	0	1	0	1	0	1	0	EFX	efdcfuf	
efdcfui	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	0	0	0	0	0	0	0	EFX	efdcfui	
efdcmpcq	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	0	1	1	1	1	0	1	1	1	0	EFX	efdcmpcq
efdcmpgt	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	0	1	1	1	0	0	1	1	0	0	EFX	efdcmpgt
efdcmpit	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	0	1	1	1	0	1	1	0	1	1	EFX	efdcmpit
efdcstf	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	1	1	1	1	1	1	1	EFX	efdcstf	
efdcstsi	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	1	0	1	0	1	1	1	EFX	efdcstsi	
efdcstsz	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	1	1	0	1	0	1	0	1	EFX	efdcstsz	
efdcstuf	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	1	1	1	0	1	1	0	EFX	efdcstuf	
efdcstui	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	0	1	0	0	0	0	0	0	EFX	efdcstui	
efdcstuiZ	0	0	0	1	0	0	rD			///			rB			0	1	0	1	1	1	1	1	1	0	0	0	0	0	0	0	EFX	efdcstuiZ	
efddiv	0	0	0	1	0	0	rD			rA			rB			0	1	0	1	1	1	1	0	1	0	0	0	1	1	0	0	EFX	efddiv	
efdmul	0	0	0	1	0	0	rD			rA			rB			0	1	0	1	1	1	1	0	1	0	0	0	0	0	0	0	EFX	efdmul	
efdnabs	0	0	0	1	0	0	rD			rA			///			0	1	0	1	1	1	1	0	0	1	0	1	0	1	1	EFX	efdnabs		
efdneg	0	0	0	1	0	0	rD			rA			///			0	1	0	1	1	1	1	0	0	1	1	0	1	1	0	1	EFX	efdneg	
efdsb	0	0	0	1	0	0	rD			rA			rB			0	1	0	1	1	1	1	0	0	0	0	0	1	1	1	1	EFX	efdsb	
efdsteq	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	1	1	1	1	1	1	1	1	0	EFX	efdsteq	
efdstgt	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	1	1	1	1	1	1	1	0	0	EFX	efdstgt	
efdstlt	0	0	0	1	0	0	crfD	/	/	rA			rB			0	1	0	1	1	1	1	1	1	1	1	1	1	0	1	1	EFX	efdstlt	



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
efsabs	0	0	0	1	0	0	rD			rA						///			0	1	0	1	1	0	0	0	1	0	0	EFX	efsabs			
efsadd	0	0	0	1	0	0	rD			rA						rB			0	1	0	1	1	0	0	0	0	0	0	0	0	0	EFX	efsadd
efscfd	0	0	0	1	0	0	rD			0	0	0	0	0	rB			0	1	0	1	1	0	0	1	1	1	1	1	1	1	1	EFX	efscfd
efscfsf	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	0	1	1	1	1	EFX	efscfsf	
efscfsi	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	0	0	0	1	1	EFX	efscfsi	
efscfuf	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	0	1	0	1	0	EFX	efscfuf	
efscfui	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	0	0	0	0	0	EFX	efscfui	
efscmpeq	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	0	1	1	1	1	0	EFX	efscmpeq		
efscmpgt	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	0	1	1	0	0	0	EFX	efscmpgt		
efscmplt	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	0	1	1	0	1	0	EFX	efscmplt		
efscstf	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	1	1	1	1	1	EFX	efscstf	
efscstsi	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	1	0	1	0	EFX	efscstsi		
efscstsz	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	1	0	1	0	1	0	EFX	efscstsz	
efscstuf	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	1	1	0	1	0	EFX	efscstuf	
efscstui	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	0	1	0	0	0	0	EFX	efscstui	
efscstuiZ	0	0	0	1	0	0	rD			///						rB			0	1	0	1	1	0	1	1	0	0	0	0	0	EFX	efscstuiZ	
efsddiv	0	0	0	1	0	0	rD			rA						rB			0	1	0	1	1	0	0	1	0	0	1	0	EFX	efsddiv		
efsmul	0	0	0	1	0	0	rD			rA						rB			0	1	0	1	1	0	0	1	0	0	0	0	EFX	efsmul		
efsnabs	0	0	0	1	0	0	rD			rA						///			0	1	0	1	1	0	0	0	1	0	1	0	EFX	efsnabs		
efsneg	0	0	0	1	0	0	rD			rA						///			0	1	0	1	1	0	0	0	1	1	0	1	0	EFX	efsneg	
efssub	0	0	0	1	0	0	rD			rA						rB			0	1	0	1	1	0	0	0	0	0	0	1	0	EFX	efssub	
efststseq	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	1	1	1	1	1	0	EFX	efststseq		
efststgt	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	1	1	1	1	0	0	EFX	efststgt		
efststlt	0	0	0	1	0	0	crfD	/	/	rA						rB			0	1	0	1	1	0	1	1	1	1	0	1	EFX	efststlt		
mulli	0	0	0	1	1	1	rD			rA						SIMM											D	mulli						
subfic	0	0	1	0	0	0	rD			rA						SIMM											D	subfic						
cmpli	0	0	1	0	1	0	crfD	/	L	rA						UIMM											D	cmpli						
cmpi	0	0	1	0	1	1	crfD	/	L	rA						SIMM											D	cmpi						



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic							
addc	0	0	1	1	0	0	rD			rA			SIMM															D	addc												
addc.	0	0	1	1	0	1	rD			rA			SIMM															D	addc.												
addi	0	0	1	1	1	0	rD			rA			SIMM															D	addi												
addis	0	0	1	1	1	1	rD			rA			SIMM															D	addis												
bc	0	1	0	0	0	0	BO			BI			BD												0	0	B	bc													
bca	0	1	0	0	0	0	BO			BI			BD												1	0	B	bca													
bcl	0	1	0	0	0	0	BO			BI			BD												0	1	B	bcl													
bcla	0	1	0	0	0	0	BO			BI			BD												1	1	B	bcla													
sc	0	1	0	0	0	1	///															1	/	SC	sc																
b	0	1	0	0	1	0	LI															0	0	I	b																
ba	0	1	0	0	1	0	LI															1	0	I	ba																
bl	0	1	0	0	1	0	LI															0	1	I	bl																
bla	0	1	0	0	1	0	LI															1	1	I	bla																
rfci	0	1	0	0	1	1	///										0	0	0	0	1	1	0	0	1	1	/	XL	rfci												
rfmci <sup>(1)</sup>	0	1	0	0	1	1	///										0	0	0	0	1	0	0	1	1	0	/	XL	rfmci												
mcrf	0	1	0	0	1	1	crfD	//	crfS	///										0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	/	XL	mcrf				
bclr	0	1	0	0	1	1	BO			BI			///												0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	XL	bclr
bclrl	0	1	0	0	1	1	BO			BI			///												0	0	0	0	0	1	0	0	0	0	0	0	1	XL	bclrl		
crnor	0	1	0	0	1	1	crbD			crbA			crbB												0	0	0	0	1	0	0	0	0	0	1	/	XL	crnor			
rfi <sup>(1)</sup>	0	1	0	0	1	1	///										0	0	0	0	1	1	0	0	1	0	/	XL	rfi												
crandc	0	1	0	0	1	1	crbD			crbA			crbB												0	0	1	0	0	0	0	0	0	0	1	/	XL	crandc			
isync	0	1	0	0	1	1	///										0	0	1	0	0	1	0	1	1	0	/	XL	isync												
crxor	0	1	0	0	1	1	crbD			crbA			crbB												0	0	1	1	0	0	0	0	0	0	1	/	XL	crxor			
crand	0	1	0	0	1	1	crbD			crbA			crbB												0	1	0	0	0	0	0	0	0	0	1	/	XL	crand			
crnand	0	1	0	0	1	1	crbD			crbA			crbB												0	0	1	1	1	0	0	0	0	0	1	/	XL	crnand			
creqv	0	1	0	0	1	1	crbD			crbA			crbB												0	1	0	0	1	0	0	0	0	0	1	/	XL	creqv			
crorc	0	1	0	0	1	1	crbD			crbA			crbB												0	1	1	0	1	0	0	0	0	0	1	/	XL	crorc			
cror	0	1	0	0	1	1	crbD			crbA			crbB												0	1	1	1	0	0	0	0	0	0	1	/	XL	cror			



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
bcctr	0	1	0	0	1	1	BO			BI			///			1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	XL	bcctr
bcctrl	0	1	0	0	1	1	BO			BI			///			1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	1	XL	bcctrl
rlwimi	0	1	0	1	0	0	rS			rA			SH			MB			ME						0	M	rlwimi							
rlwimi.	0	1	0	1	0	0	rS			rA			SH			MB			ME						1	M	rlwimi.							
rlwinm	0	1	0	1	0	1	rS			rA			SH			MB			ME						0	M	rlwinm							
rlwinm.	0	1	0	1	0	1	rS			rA			SH			MB			ME						1	M	rlwinm.							
rlwnm	0	1	0	1	1	1	rS			rA			rB			MB			ME						0	M	rlwnm							
rlwnm.	0	1	0	1	1	1	rS			rA			rB			MB			ME						1	M	rlwnm.							
ori	0	1	1	0	0	0	rS			rA			UIMM															D	ori					
oris	0	1	1	0	0	1	rS			rA			UIMM															D	oris					
xori	0	1	1	0	1	0	rS			rA			UIMM															D	xori					
xoris	0	1	1	0	1	1	rS			rA			UIMM															D	xoris					
andi.	0	1	1	1	0	0	rS			rA			UIMM															D	andi.					
andis.	0	1	1	1	0	1	rS			rA			UIMM															D	andis.					
dcblc	0	1	1	1	1	1	CT			rA			rB			0	1	1	0	0	0	0	0	0	1	1	0	0			X	dcblc		
dcbtIs	0	1	1	1	1	1	CT			rA			rB			0	0	1	0	1	0	0	0	0	1	1	0	0			X	dcbtIs		
dcbtstIs	0	1	1	1	1	1	CT			rA			rB			0	0	1	0	0	0	0	0	0	1	1	0	0			X	dcbtstIs		
evabs	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	1	0	0	0			EVX	evabs		
evaddiw	0	1	1	1	1	1	rD			UIMM			rB			0	1	0	0	0	0	0	0	0	0	0	1	0			EVX	evaddiw		
evaddsmiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	1	0	0	0	1			EVX	evaddsmiaaw		
evaddssiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	0	1			EVX	evaddssiaaw	
evaddumiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	1	0	0	0	0			EVX	evaddumiaaw		
evaddusiaaw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	0			EVX	evaddusiaaw		
evaddw	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	0	0	0	0	0	0	0			EVX	evaddw	
evand	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	0	0	0	1			EVX	evand		
evandc	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	0	1	0			EVX	evandc			



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evcmpeq	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	0	0	1	1	0	1	0	0	EVX	evcmpeq			
evcmpgts	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	0	0	1	1	0	0	0	1	EVX	evcmpgts			
evcmpgtu	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	0	0	1	1	0	0	0	0	EVX	evcmpgtu			
evcmplts	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	0	0	1	1	0	0	1	1	EVX	evcmplts			
evcmpltu	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	0	0	1	1	0	0	1	0	EVX	evcmpltu			
evcntlsw	0	1	1	1	1	1	rD				rA				///				0	1	0	0	0	0	0	1	1	1	0	EVX	evcntlsw			
evcntlzw	0	1	1	1	1	1	rD				rA				///				0	1	0	0	0	0	0	1	1	0	1	EVX	evcntlzw			
evdivws	0	1	1	1	1	1	rD				rA				rB				1	0	0	1	1	0	0	0	1	1	0	EVX	evdivws			
evdivwu	0	1	1	1	1	1	rD				rA				rB				1	0	0	1	1	0	0	0	1	1	1	EVX	evdivwu			
eveqv	0	1	1	1	1	1	rD				rA				rB				0	1	0	0	0	0	1	1	0	0	1	EVX	eveqv			
evextsb	0	1	1	1	1	1	rD				rA				///				0	1	0	0	0	0	0	1	0	1	0	EVX	evextsb			
evextsh	0	1	1	1	1	1	rD				rA				///				0	1	0	0	0	0	0	1	0	1	1	EVX	evextsh			
evfsabs	0	1	1	1	1	1	rD				rA				///				0	1	0	1	0	0	0	0	1	0	0	EVX	evfsabs			
evfsadd	0	1	1	1	1	1	rD				rA				rB				0	1	0	1	0	0	0	0	0	0	0	EVX	evfsadd			
evfscfsf	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	0	1	1	EVX	evfscfsf			
evfscfsi	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	0	0	1	EVX	evfscfsi			
evfscfuf	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	0	1	0	EVX	evfscfuf			
evfscfui	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	0	0	0	EVX	evfscfui			
evfscmpeq	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	1	0	0	0	1	1	1	0	EVX	evfscmpeq			
evfscmpgt	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	1	0	0	0	1	1	0	0	EVX	evfscmpgt			
evfscmplt	0	1	1	1	1	1	crfD		/	/	rA				rB				0	1	0	1	0	0	0	1	1	0	1	EVX	evfscmplt			
evfsctsf	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	1	1	1	EVX	evfsctsf			
evfsctsi	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	1	0	1	EVX	evfsctsi			
evfsctsiz	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	1	0	1	0	EVX	evfsctsiz			
evfsctuf	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	1	1	0	EVX	evfsctuf			
evfsctui	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	0	1	0	0	EVX	evfsctui			
evfsctuiz	0	1	1	1	1	1	rD				///				rB				0	1	0	1	0	0	1	1	0	0	0	EVX	evfsctuiz			
evfsdiv	0	1	1	1	1	1	rD				rA				rB				0	1	0	1	0	0	0	1	0	0	1	EVX	evfsdiv			



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evfsmul	0	1	1	1	1	1	rD			rA			rB			0	1	0	1	0	0	0	1	0	0	0	1	0	0	0	EVX	evfsmul		
evfsnabs	0	1	1	1	1	1	rD			rA			///			0	1	0	1	0	0	0	0	0	1	0	1	0	1	0	EVX	evfsnabs		
evfsneg	0	1	1	1	1	1	rD			rA			///			0	1	0	1	0	0	0	0	0	1	1	0	1	1	0	EVX	evfsneg		
evfssub	0	1	1	1	1	1	rD			rA			rB			0	1	0	1	0	0	0	0	0	0	0	0	0	1	0	EVX	evfssub		
evfststeq	0	1	1	1	1	1	crfD		/	/	rA			rB			0	1	0	1	0	0	1	1	1	1	1	1	0	EVX	evfststeq			
evfststgt	0	1	1	1	1	1	crfD		/	/	rA			rB			0	1	0	1	0	0	1	1	1	1	0	0	0	EVX	evfststgt			
evfstslt	0	1	1	1	1	1	crfD		/	/	rA			rB			0	1	0	1	0	0	1	1	1	1	0	1	0	EVX	evfstslt			
evladd	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	0	0	0	0	0	0	1	EVX	evladd			
evladdx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evladdx			
evldh	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	0	0	1	0	1	0	1	EVX	evldh			
evldhx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	0	0	1	0	0	0	EVX	evldhx				
evldw	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	0	0	0	1	1	0	EVX	evldw				
evldwx	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	0	0	0	1	0	0	EVX	evldwx				
evlhhesplat	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	1	0	0	0	1	0	EVX	evlhhesplat				
evlhhesplat <sub>x</sub>	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	1	0	0	0	0	0	EVX	evlhhesplat <sub>x</sub>				
evlhhossp <sub>at</sub>	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	1	1	1	1	0	EVX	evlhhossp <sub>at</sub>					
evlhhossp <sub>atx</sub>	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	1	1	1	0	0	EVX	evlhhossp <sub>atx</sub>					
evlhhou <sub>spat</sub>	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	0	1	1	0	1	0	EVX	evlhhou <sub>spat</sub>					
evlhhou <sub>spatx</sub>	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	0	1	1	0	0	0	EVX	evlhhou <sub>spatx</sub>					
evlwhe	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	1	0	0	0	1	0	1	EVX	evlwhe				
evlwhex	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	0	0	0	0	0	EVX	evlwhex				
evlw hos	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	1	0	1	1	1	0	EVX	evlw hos					
evlw hos <sub>x</sub>	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	1	1	0	0	EVX	evlw hos <sub>x</sub>					
evlw hou	0	1	1	1	1	1	rD			rA			UIMM <sup>(2)</sup>			0	1	1	0	0	0	1	0	1	0	1	0	EVX	evlw hou					
evlw hou <sub>x</sub>	0	1	1	1	1	1	rD			rA			rB			0	1	1	0	0	0	1	0	1	0	0	0	EVX	evlw hou <sub>x</sub>					





Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evlwhsplat	0	1	1	1	1	1																0	1	1	0	0	0	1	1	1	0	1	EVX	evlwhsplat
evlwhsplat x	0	1	1	1	1	1																0	1	1	0	0	0	1	1	1	0	0	EVX	evlwhsplatx
evlwwsplat	0	1	1	1	1	1																0	1	1	0	0	0	1	1	0	0	1	EVX	evlwwsplat
evlwwsplat x	0	1	1	1	1	1																0	1	1	0	0	0	1	1	0	0	0	EVX	evlwwsplatx
evmergehi	0	1	1	1	1	1																0	1	0	0	0	1	0	1	1	0	0	EVX	evmergehi
evmergehil o	0	1	1	1	1	1																0	1	0	0	0	1	0	1	1	1	0	EVX	evmergehilo
evmergeelo	0	1	1	1	1	1																0	1	0	0	0	1	0	1	1	0	1	EVX	evmergeelo
evmergeelo hi	0	1	1	1	1	1																0	1	0	0	0	1	0	1	1	1	1	EVX	evmergeelohi
evmhegsmf aa	0	1	1	1	1	1																1	0	1	0	0	1	0	1	0	1	1	EVX	evmhegsmf aa
evmhegsmf an	0	1	1	1	1	1																1	0	1	1	0	1	0	1	0	1	1	EVX	evmhegsmf an
evmhegsmi aa	0	1	1	1	1	1																1	0	1	0	0	1	0	1	0	0	1	EVX	evmhegsmi aa
evmhegsmi an	0	1	1	1	1	1																1	0	1	1	0	1	0	1	0	0	1	EVX	evmhegsmi an
evmhegumi aa	0	1	1	1	1	1																1	0	1	0	0	1	0	1	0	0	0	EVX	evmhegumi aa
evmhegumi an	0	1	1	1	1	1																1	0	1	1	0	1	0	1	0	0	0	EVX	evmhegumi an
evmhesmf	0	1	1	1	1	1																1	0	0	0	0	0	0	1	0	1	1	EVX	evmhesmf
evmhesmf a	0	1	1	1	1	1																1	0	0	0	0	1	0	1	0	1	1	EVX	evmhesmf a
evmhesmf aw	0	1	1	1	1	1																1	0	1	0	0	0	0	1	0	1	1	EVX	evmhesmf aw
evmhesmf nw	0	1	1	1	1	1																1	0	1	1	0	0	0	1	0	1	1	EVX	evmhesmf nw
evmhesmi	0	1	1	1	1	1																1	0	0	0	0	0	0	1	0	0	1	EVX	evmhesmi
evmhesmia	0	1	1	1	1	1																1	0	0	0	0	1	0	1	0	0	1	EVX	evmhesmia
evmhesmia aw	0	1	1	1	1	1																1	0	1	0	0	0	0	1	0	0	1	EVX	evmhesmia aw



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhesmia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	0	0	1	0	0	1	EVX	evmhesmia nw			
evmhessf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	EVX	evmhessf		
evmhessfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	1	0	0	0	0	1	1	EVX	evmhessfa				
evmhessfa aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	1	1	EVX	evmhessfaa w			
evmhessfa nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	1	1	EVX	evmhessfan w			
evmhessia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	0	1	EVX	evmhessiaa w			
evmhessia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	0	1	EVX	evmhessian w			
evmheumi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	0	0	0	0	0	EVX	evmheumi			
evmheumia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	1	0	1	0	0	0	0	0	EVX	evmheumia			
evmheumia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	0	0	0	0	0	EVX	evmheumia aw			
evmheumia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	1	0	0	0	0	0	EVX	evmheumia nw			
evmheusia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	0	0	0	0	0	0	EVX	evmheusiaa w		
evmheusia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evmheusian w		
evmhogsm faa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	1	1	1	1	1	EVX	evmhogsmf aa			
evmhogsm fan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	1	1	1	1	1	EVX	evmhogsmf an			
evmhogsmi aa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	1	0	1	EVX	evmhogsmi aa					
evmhogsmi an	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	1	0	1	EVX	evmhogsmi an					
evmhogum iaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	1	0	1	1	1	0	0	EVX	evmhogumi aa					
evmhogum ian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	1	0	1	1	1	0	0	EVX	evmhogumi an					
evmhosmf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	0	1	1	1	1	EVX	evmhosmf				



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmhosmfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	1	1	1	1	1	1	1	EVX	evmhosmfa		
evmhosmfa aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	1	1	1	1	1	1	1	1	1	EVX	evmhosmfa aw	
evmhosmfa nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	1	1	1	1	1	1	1	EVX	evmhosmfa nw		
evmhosmi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	1	1	0	1	1	1	1	1	EVX	evmhosmi		
evmhosmia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	1	0	1	1	1	1	1	EVX	evmhosmia		
evmhosmia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	1	1	0	1	1	1	1	1	EVX	evmhosmia aw		
evmhosmia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	1	0	1	1	1	1	1	EVX	evmhosmia nw		
evmhossf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	0	1	1	1	1	1	1	1	EVX	evmhossf		
evmhossfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	0	1	1	1	1	1	1	1	EVX	evmhossfa		
evmhossfa aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	1	1	1	1	1	1	EVX	evmhossfaa w		
evmhossfa nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	1	1	1	1	1	1	EVX	evmhossfan w		
evmhossia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	0	1	0	1	1	1	1	EVX	evmhossiaa w		
evmhossia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	0	1	0	1	1	1	1	EVX	evmhossian w		
evmhoumi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	0	0	1	1	0	0	1	1	1	1	EVX	evmhoumi		
evmhoumia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	0	1	0	1	1	0	0	1	1	1	1	EVX	evmhoumia		
evmhoumia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	1	1	0	0	1	1	1	1	EVX	evmhoumia aw		
evmhoumia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	1	1	0	0	1	1	1	1	EVX	evmhoumia nw		
evmhousia aw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	0	0	0	0	1	0	0	1	1	1	1	EVX	evmhousiaa w		
evmhousia nw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	0	0	0	0	1	0	0	1	1	1	1	EVX	evmhousian w		
evmra	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	1	0	0	1	1	1	1	EVX	evmra		
evmwhgsm faa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	1	1	0	1	1	1	1	1	1	1	EVX	evmwhgsmf aa		



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmwhgsmfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	1	1	1	1	1	1	1	1	EVX	evmwhgsmfan	
evmwhgsmiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	1	0	1	1	1	0	1	1	0	1	1	EVX	evmwhgsmiaa	
evmwhgsmian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	1	1	0	1	1	EVX	evmwhgsmian	
evmwhgssfaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	1	0	0	1	1	0	0	1	1	1	1	EVX	evmwhgssfaa	
evmwhgssfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	0	1	1	1	1	1	1	1	1	EVX	evmwhgssfan	
evmwhgumiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	1	0	1	1	1	0	1	1	0	0	0	EVX	evmwhgumiaa	
evmwhgumian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	1	0	1	1	0	0	0	EVX	evmwhgumian		
evmwhsmf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmf	
evmwhsmfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfa	
evmwhsmfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfaaw	
evmwhsmfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	0	1	1	1	1	1	1	1	1	1	EVX	evmwhsmfanw	
evmwhsmi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	0	1	1	0	1	1	0	1	1	EVX	evmwhsmi		
evmwhsmia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	0	1	1	0	1	1	0	1	1	EVX	evmwhsmia		
evmwhsmiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	0	1	1	0	1	1	0	1	1	EVX	evmwhsmiaaw		
evmwhsmianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	0	1	1	0	1	1	0	1	1	EVX	evmwhsmianw		
evmwhssf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	0	0	1	1	1	1	1	1	1	1	EVX	evmwhssf	
evmwhssfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	0	0	1	1	1	1	1	1	1	1	EVX	evmwhssfa	
evmwhssfaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	0	0	1	1	1	1	1	1	1	1	EVX	evmwhssfaaw	
evmwhssfanw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	0	0	1	1	1	1	1	1	1	1	EVX	evmwhssfanw	
evmwhssianw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	0	0	1	0	1	0	1	0	1	1	EVX	evmwhssianw	



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmwhssm aaw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	0	1	0	0	0	1	0	1	0	1	EVX	evmwhssm aaw			
evmwhumi	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	0	0	1	1	0	0	1	1	0	0	EVX	evmwhumi					
evmwhumi a	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	1	0	1	1	0	0	1	1	0	0	EVX	evmwhumia					
evmwhusia aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	0	1	0	0	0	1	0	0	EVX	evmwhusia aw					
evmwhusia nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	0	1	0	0	0	1	0	0	EVX	evmwhusia nw					
evmwlsmf	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	0	0	1	0	1	1	0	1	1	EVX	evmwlsmf						
evmwlsmf a	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	1	0	1	0	1	1	0	1	1	EVX	evmwlsmf a						
evmwlsmf aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	1	0	1	1	0	1	1	EVX	evmwlsmf aw						
evmwlsmf nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	1	0	1	1	0	1	1	EVX	evmwlsmf nw						
evmwlsmia aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	1	0	0	1	0	0	1	EVX	evmwlsmia aw						
evmwlsmia nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	1	0	0	1	0	0	1	EVX	evmwlsmia nw						
evmwlsf	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	0	0	0	0	1	1	0	0	0	1	1	EVX	evmwlsf				
evmwlsf a	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	1	0	0	0	1	1	0	0	0	1	1	EVX	evmwlsf a				
evmwlsf aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	0	0	1	1	0	0	0	0	1	1	EVX	evmwlsf aw			
evmwlsf nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	0	0	1	1	0	0	0	0	1	1	EVX	evmwlsf nw			
evmwlssia aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	0	0	0	1	0	0	0	0	1	EVX	evmwlssia aw				
evmwlssia nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	0	0	0	1	0	0	0	0	1	EVX	evmwlssia nw				
evmwlumi	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	0	0	1	0	0	0	1	0	0	0	EVX	evmwlumi					
evmwlumia	0	1	1	1	1	1	rD		rA		rB		1	0	0	0	1	1	0	1	0	0	0	0	0	0	0	EVX	evmwlumia					
evmwlumia aw	0	1	1	1	1	1	rD		rA		rB		1	0	1	0	1	0	0	1	0	0	1	0	0	0	0	EVX	evmwlumia aw					
evmwlumia nw	0	1	1	1	1	1	rD		rA		rB		1	0	1	1	1	0	0	1	0	0	1	0	0	0	0	EVX	evmwlumia nw					



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evmwlusiaaw	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evmwlusiaaw	
evmwlusiannw	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	EVX	evmwlusiannw	
evmwsmf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	1	0	1	1	0	1	1	0	0	EVX	evmwsmf	
evmwsmfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	1	1	0	1	1	0	1	1	0	0	EVX	evmwsmfa	
evmwsmfaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	1	1	0	1	1	0	1	1	0	0	EVX	evmwsmfaa	
evmwsmfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	0	1	1	0	1	1	0	0	EVX	evmwsmfan	
evmwsmi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	1	0	1	1	0	0	1	0	0	EVX	evmwsmi	
evmwsmia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	1	1	0	0	1	0	0	1	0	0	EVX	evmwsmia	
evmwsmiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	1	1	0	1	1	0	0	1	0	0	EVX	evmwsmiaa	
evmwsmian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	0	1	1	0	0	1	0	0	EVX	evmwsmian	
evmwssf	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	0	0	1	0	0	1	1	0	0	EVX	evmwssf	
evmwssfa	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	1	0	0	1	1	0	0	1	1	0	0	EVX	evmwssfa
evmwssfaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	1	0	1	0	0	1	1	0	0	1	1	EVX	evmwssfaa
evmwssfan	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	0	0	1	0	0	1	1	0	0	EVX	evmwssfan	
evmwumi	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	0	1	1	0	0	0	0	0	0	0	0	0	EVX	evmwumi
evmwumia	0	1	1	1	1	1	rD			rA			rB			1	0	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	EVX	evmwumia
evmwumiaa	0	1	1	1	1	1	rD			rA			rB			1	0	1	0	1	0	1	1	0	1	1	0	0	0	0	0	0	EVX	evmwumiaa
evmwumian	0	1	1	1	1	1	rD			rA			rB			1	0	1	1	1	0	1	1	0	1	1	0	0	0	0	0	0	EVX	evmwumian
evnand	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	EVX	evnand	
evneg	0	1	1	1	1	1	rD			rA			///			0	1	0	0	0	0	0	0	0	0	1	0	0	1	0	0	EVX	evneg	
evnor	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	EVX	evnor
evor	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	1	0	1	1	1	0	0	EVX	evor	
evorc	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	1	0	1	1	0	1	1	0	0	EVX	evorc
evrlw	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	EVX	evrlw



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evrlwi	0	1	1	1	1	1			rD		rA			UIMM			0	1	0	0	0	1	0	1	0	1	0	1	0		EVX	evrlwi		
evrndw	0	1	1	1	1	1			rD		rA			UIMM			0	1	0	0	0	0	0	0	0	0	1	1	0	0		EVX	evrndw	
evsel	0	1	1	1	1	1			rD		rA			rB			0	1	0	0	1	1	1	1					crfS		EVX	evsel		
evslw	0	1	1	1	1	1			rD		rA			rB			0	1	0	0	0	1	0	0	0	1	0	0	1	0	0	EVX	evslw	
evslwi	0	1	1	1	1	1			rD		rA			UIMM			0	1	0	0	0	1	0	0	1	0	0	1	1	0		EVX	evslwi	
evsplatfi	0	1	1	1	1	1			rD		SIMM			///			0	1	0	0	0	1	0	1	0	1	0	1	1		EVX	evsplatfi		
evsplati	0	1	1	1	1	1			rD		SIMM			///			0	1	0	0	0	1	0	1	0	0	1	0	0	1		EVX	evsplati	
evsrwis	0	1	1	1	1	1			rD		rA			UIMM			0	1	0	0	0	1	0	0	0	0	1	1			EVX	evsrwis		
evsrwiu	0	1	1	1	1	1			rD		rA			UIMM			0	1	0	0	0	1	0	0	0	0	1	0	0	1	0		EVX	evsrwiu
evsrws	0	1	1	1	1	1			rD		rA			rB			0	1	0	0	0	1	0	0	0	0	0	0	0	1		EVX	evsrws	
evsrwu	0	1	1	1	1	1			rD		rA			rB			0	1	0	0	0	1	0	0	0	0	0	0	0	0		EVX	evsrwu	
evstdd	0	1	1	1	1	1			rD		rA			UIMM <sup>(2)</sup>			0	1	1	0	0	1	0	0	0	0	0	0	0	1		EVX	evstdd	
evstddx	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	0	0	0	0	0	0	0		EVX	evstddx		
evstdh	0	1	1	1	1	1			rS		rA			UIMM <sup>(2)</sup>			0	1	1	0	0	1	0	0	1	0	0	1	0	1		EVX	evstdh	
evstdhx	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	0	0	1	0	0	1	0	0		EVX	evstdhx	
evstdw	0	1	1	1	1	1			rS		rA			UIMM <sup>(2)</sup>			0	1	1	0	0	1	0	0	0	0	0	1	1		EVX	evstdw		
evstdwx	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	0	0	0	0	1	0			EVX	evstdwx		
evstwhe	0	1	1	1	1	1			rS		rA			UIMM <sup>(3)</sup>			0	1	1	0	0	1	1	0	0	0	0	0	1		EVX	evstwhe		
evstwhex	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	1	0	0	0	0	0	0		EVX	evstwhex		
evstwho	0	1	1	1	1	1			rS		rA			UIMM <sup>(3)</sup>			0	1	1	0	0	1	1	0	1	0	1	0	1		EVX	evstwho		
evstwhox	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	1	0	1	0	0	0		EVX	evstwhox			
evstwwe	0	1	1	1	1	1			rS		rA			UIMM <sup>(3)</sup>			0	1	1	0	0	1	1	1	1	0	0	1		EVX	evstwwe			
evstwwex	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	1	1	0	0	0	0		EVX	evstwwex			
evstwwo	0	1	1	1	1	1			rS		rA			UIMM <sup>(3)</sup>			0	1	1	0	0	1	1	1	1	1	0	1		EVX	evstwwo			
evstwwox	0	1	1	1	1	1			rS		rA			rB			0	1	1	0	0	1	1	1	1	1	0	0		EVX	evstwwox			
evsubsmia aw	0	1	1	1	1	1			rD		rA			///			1	0	0	1	1	0	0	1	0	0	1	0	1	1		EVX	evsubsmia aw	



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
evsubfssia aw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	0	1	1	EVX	evsubfssia w		
evsubfumia aw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	1	0	0	1	0	1	0	EVX	evsubfumia aw			
evsubfusia aw	0	1	1	1	1	1	rD			rA			///			1	0	0	1	1	0	0	0	0	0	0	0	1	0	EVX	evsubfusia w			
evsubfw	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	0	0	0	0	1	0	0	EVX	evsubfw			
evsubifw	0	1	1	1	1	1	rD			UIMM			rB			0	1	0	0	0	0	0	0	0	0	1	1	0	EVX	evsubifw				
evxor	0	1	1	1	1	1	rD			rA			rB			0	1	0	0	0	0	0	1	0	1	1	0	EVX	evxor					
icblc	0	1	1	1	1	1	CT			rA			rB			0	0	1	1	1	0	0	1	1	0	0	X	icblc						
icbt	0	1	1	1	1	1	CT			rA			rB			0	0	0	0	0	1	0	1	1	0	/	X	icbt						
icbtls	0	1	1	1	1	1	CT			rA			rB			0	1	1	1	1	0	0	1	1	0	0	X	icbtls						
isel	0	1	1	1	1	1	rD			rA			rB			crb			0	1	1	1	1	0	X	isel								
mbar	0	1	1	1	1	1	MO			///			1	1	0	1	0	1	0	1	1	0	/	X	mbar									
mfdc	0	1	1	1	1	1	rD			DCRN5-9			DCRN0-4			0	1	0	1	0	0	0	0	1	1	/	XFX	mfdc						
mfpmr	0	1	1	1	1	1	rD			PMRN5-9			PMRN0-4			0	1	0	1	0	0	1	1	1	0	0	XFX	mfpmr						
msync	0	1	1	1	1	1	///			///			1	0	0	1	0	1	0	1	1	0	/	X	msync									
mtdc	0	1	1	1	1	1	rS			DCRN5-9			DCRN0-4			0	1	1	1	0	0	0	0	1	1	/	XFX	mtdc						
mtpmr	0	1	1	1	1	1	rS			PMRN5-9			PMRN0-4			0	1	1	1	0	0	1	1	1	0	0	XFX	mtpmr						
tlbivax	0	1	1	1	1	1	///			rA			rB			1	1	0	0	0	1	0	0	1	0	/	X	tlbivax						
tlbre	0	1	1	1	1	1	/// <sup>(4)</sup>			///			1	1	1	0	1	1	0	0	1	0	/	X	tlbre									
tlbsx	0	1	1	1	1	1	/// <sup>(5)</sup>			rA			rB			1	1	1	0	0	1	0	0	1	0	/	X	tlbsx						
tlbwe	0	1	1	1	1	1	/// <sup>(5)</sup>			///			1	1	1	1	0	1	0	0	1	0	/	X	tlbwe									
wait	0	1	1	1	1	1	///			///			0	0	0	0	1	1	1	1	1	0	/	X	wait									
wrtree	0	1	1	1	1	1	rS			///			0	0	1	0	0	0	0	0	0	0	1	1	/	X	wrtree							
wrtreei	0	1	1	1	1	1	///			E	///			0	0	1	0	1	0	0	0	0	1	1	/	X	wrtreei							
cmp	0	1	1	1	1	1	crfD		/	L	rA			rB			0	0	0	0	0	0	0	0	0	/	X	cmp						
tw	0	1	1	1	1	1	TO			rA			rB			0	0	0	0	0	0	0	1	0	/	X	tw							
subfc	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	1	0	0	0	X	subfc							





Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic			
subfc.	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1	X	subfc.			
addc	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	X	addc			
addc.	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0	1	X	addc.			
mulhwu	0	1	1	1	1	1	rD			rA			rB			/	0	0	0	0	0	0	0	0	0	0	1	0	1	1	0	X	mulhwu				
mulhwu.	0	1	1	1	1	1	rD			rA			rB			/	0	0	0	0	0	0	0	0	0	0	1	0	1	1	1	X	mulhwu.				
mfcrr	0	1	1	1	1	1	rD			///			///			0	0	0	0	0	0	0	0	0	0	1	0	0	1	1	/	X	mfcrr				
lwarx	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	/	X	lwarx				
lwzxx	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	1	0	1	1	1	/	X	lwzxx				
slw	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	X	slw				
slw.	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	1	X	slw.			
cntlzw	0	1	1	1	1	1	rS			rA			///			0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	0	X	cntlzw				
cntlzw.	0	1	1	1	1	1	rS			rA			///			0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	1	X	cntlzw.				
and	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	1	0	0	0	X	and				
and.	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	1	0	0	0	1	X	and.			
cmpl	0	1	1	1	1	1	/	L	rA			rB			///			0	0	0	0	0	0	0	0	0	0	0	0	0	0	/	X	cmpl			
subf	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	X	subf			
subf.	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	1	X	subf.		
dcbst	0	1	1	1	1	1	///			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	0	1	1	0	/	X	dcbst			
lwzux	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	0	1	1	0	1	1	1	/	X	lwzux		
andc	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	1	1	0	0	0	X	andc			
andc.	0	1	1	1	1	1	rS			rA			rB			0	0	0	0	0	0	0	0	0	0	1	1	1	1	0	0	0	1	X	andc.		
mulhw	0	1	1	1	1	1	rD			rA			rB			/	0	0	0	0	0	0	0	0	0	0	1	0	1	1	0	X	mulhw				
mulhw.	0	1	1	1	1	1	rD			rA			rB			/	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	1	1	X	mulhw.		
mfmsr <sup>(1)</sup>	0	1	1	1	1	1	rD			///			///			0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	1	/	X	mfmsr		
dcbf	0	1	1	1	1	1	///			rA			rB			0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	1	0	/	X	dcbf	
lbzx	0	1	1	1	1	1	rD			rA			rB			0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1	1	/	X	lbzx		
neg	0	1	1	1	1	1	rD			rA			///			0	0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	0	0	X	neg		
neg.	0	1	1	1	1	1	rD			rA			///			0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	1	0	0	0	1	X	neg.



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
lbzux	0	1	1	1	1	1			rD					rA					rB			0	0	0	1	1	1	0	1	1	1	/	X	lbzux
nor	0	1	1	1	1	1			rS					rA					rB			0	0	0	1	1	1	1	1	0	0	0	X	nor
nor.	0	1	1	1	1	1			rS					rA					rB			0	0	0	1	1	1	1	1	0	0	1	X	nor.
subfe	0	1	1	1	1	1			rD					rA					rB			0	0	1	0	0	0	1	0	0	0	0	X	subfe
subfe.	0	1	1	1	1	1			rD					rA					rB			0	0	1	0	0	0	1	0	0	0	1	X	subfe.
adde	0	1	1	1	1	1			rD					rA					rB			0	0	1	0	0	0	1	0	1	0	0	X	adde
adde.	0	1	1	1	1	1			rD					rA					rB			0	0	1	0	0	0	1	0	1	0	1	X	adde.
mtrcf	0	1	1	1	1	1			rS	/				CRM					/			0	0	1	0	0	1	0	0	0	0	/	AFX	mtrcf
mtmsr <sup>(1)</sup>	0	1	1	1	1	1			rS					///								0	0	1	0	0	1	0	0	1	0	/	X	mtmsr
stwcx.	0	1	1	1	1	1			rS					rA					rB			0	0	1	0	0	1	0	1	1	0	1	X	stwcx.
stwx	0	1	1	1	1	1			rS					rA					rB			0	0	1	0	0	1	0	1	1	1	/	D	stwx
stwux	0	1	1	1	1	1			rS					rA					rB			0	0	1	0	1	1	0	1	1	1	/	D	stwux
subfze	0	1	1	1	1	1			rD					rA					///			0	0	1	1	0	0	1	0	0	0	0	X	subfze
subfze.	0	1	1	1	1	1			rD					rA					///			0	0	1	1	0	0	1	0	0	0	1	X	subfze.
addze	0	1	1	1	1	1			rD					rA					///			0	0	1	1	0	0	1	0	1	0	0	X	addze
addze.	0	1	1	1	1	1			rD					rA					///			0	0	1	1	0	0	1	0	1	0	1	X	addze.
stbx	0	1	1	1	1	1			rS					rA					rB			0	0	1	1	0	1	0	1	1	1	0	X	stbx
subfme	0	1	1	1	1	1			rD					rA					///			0	0	1	1	1	0	1	0	0	0	0	X	subfme
subfme.	0	1	1	1	1	1			rD					rA					///			0	0	1	1	1	0	1	0	0	0	1	X	subfme.
addme	0	1	1	1	1	1			rD					rA					///			0	0	1	1	1	0	1	0	1	0	0	X	addme
addme.	0	1	1	1	1	1			rD					rA					///			0	0	1	1	1	0	1	0	1	0	1	X	addme.
mullw	0	1	1	1	1	1			rD					rA					rB			0	0	1	1	1	0	1	0	1	1	0	X	mullw
mullw.	0	1	1	1	1	1			rD					rA					rB			0	0	1	1	1	0	1	0	1	1	1	X	mullw.
dcbtst	0	1	1	1	1	1			CT					rA					rB			0	0	1	1	1	1	0	1	1	0	/	X	dcbtst
stbux	0	1	1	1	1	1			rS					rA					rB			0	0	1	1	1	1	0	1	1	1	0	X	stbux
add	0	1	1	1	1	1			rD					rA					rB			0	1	0	0	0	0	1	0	1	0	0	X	add
add.	0	1	1	1	1	1			rD					rA					rB			0	1	0	0	0	0	1	0	1	0	1	X	add.
dcbt	0	1	1	1	1	1			CT					rA					rB			0	1	0	0	0	1	0	1	1	0	/	X	dcbt



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic				
lhzx	0	1	1	1	1	1																													X	lhzx		
eqv	0	1	1	1	1	1																														X	eqv	
eqv.	0	1	1	1	1	1																														X	eqv.	
tlbie <sup>(1),(6)</sup>	0	1	1	1	1	1																														X	tlbie	
lhzux	0	1	1	1	1	1																														X	lhzux	
xor	0	1	1	1	1	1																														X	xor	
xor.	0	1	1	1	1	1																														X	xor.	
mfspr <sup>(7)</sup>	0	1	1	1	1	1																														XFX	mfspr	
lhax	0	1	1	1	1	1																														X	lhax	
lhaux	0	1	1	1	1	1																															X	lhaux
sthx	0	1	1	1	1	1																															X	sthx
orc	0	1	1	1	1	1																															X	orc
orc.	0	1	1	1	1	1																															X	orc.
sthux	0	1	1	1	1	1																															X	sthux
or	0	1	1	1	1	1																															X	or
or.	0	1	1	1	1	1																															X	or.
divwu	0	1	1	1	1	1																															X	divwu
divwu.	0	1	1	1	1	1																															X	divwu.
mtspr <sup>(7)</sup>	0	1	1	1	1	1																															XFX	mtspr
dcbi <sup>(1)</sup>	0	1	1	1	1	1																															X	dcbi
nand	0	1	1	1	1	1																															X	nand
nand.	0	1	1	1	1	1																															X	nand.
divw	0	1	1	1	1	1																															X	divw
divw.	0	1	1	1	1	1																															X	divw.
mcrxr	0	1	1	1	1	1																															X	mcrxr
subfco	0	1	1	1	1	1																															X	subfco
subfco.	0	1	1	1	1	1																															X	subfco.
addco	0	1	1	1	1	1																															X	addco



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
addco.	0	1	1	1	1	1			rD		rA			rB			1	0	0	0	0	0	0	0	0	0	1	0	1	0	1	X	addco.	
lwbrx	0	1	1	1	1	1			rD		rA			rB			1	0	0	0	0	0	0	1	0	1	1	1	0	/	X	lwbrx		
srw	0	1	1	1	1	1			rS		rA			rB			1	0	0	0	0	0	1	1	0	0	0	0	0	0	X	srw		
srw.	0	1	1	1	1	1			rS		rA			rB			1	0	0	0	0	0	1	1	0	0	0	0	1	X	srw.			
subfo	0	1	1	1	1	1			rD		rA			rB			1	0	0	0	0	1	0	1	0	0	0	0	0	0	X	subfo		
subfo.	0	1	1	1	1	1			rD		rA			rB			1	0	0	0	0	1	0	1	0	0	0	0	0	1	X	subfo.		
tlbsync (1),(6)	0	1	1	1	1	1			///		///			///			1	0	0	0	0	1	1	0	1	1	0	1	0	/	X	tlbsync		
nego	0	1	1	1	1	1			rD		rA			///			1	0	0	1	1	0	1	0	1	0	0	0	0	0	X	nego		
nego.	0	1	1	1	1	1			rD		rA			///			1	0	0	1	1	0	1	0	1	0	0	0	1	X	nego.			
subfeo	0	1	1	1	1	1			rD		rA			rB			1	0	1	0	0	0	0	1	0	0	0	0	0	0	X	subfeo		
subfeo.	0	1	1	1	1	1			rD		rA			rB			1	0	1	0	0	0	0	1	0	0	0	0	1	X	subfeo.			
addeo	0	1	1	1	1	1			rD		rA			rB			1	0	1	0	0	0	0	1	0	1	0	1	0	0	X	addeo		
addeo.	0	1	1	1	1	1			rD		rA			rB			1	0	1	0	0	0	0	1	0	1	0	1	0	1	X	addeo.		
stwbrx	0	1	1	1	1	1			rS		rA			rB			1	0	1	0	0	1	0	1	1	1	0	/	X	stwbrx				
subfzeo	0	1	1	1	1	1			rD		rA			///			1	0	1	1	0	0	1	0	0	0	0	0	0	0	X	subfzeo		
subfzeo.	0	1	1	1	1	1			rD		rA			///			1	0	1	1	0	0	1	0	0	0	0	1	X	subfzeo.				
addzeo	0	1	1	1	1	1			rD		rA			///			1	0	1	1	0	0	1	0	1	0	1	0	0	X	addzeo			
addzeo.	0	1	1	1	1	1			rD		rA			///			1	0	1	1	0	0	1	0	1	0	1	0	1	X	addzeo.			
subfmeo	0	1	1	1	1	1			rD		rA			///			1	0	1	1	1	0	1	0	0	0	0	0	0	X	subfmeo			
subfmeo.	0	1	1	1	1	1			rD		rA			///			1	0	1	1	1	0	1	0	0	0	0	1	X	subfmeo.				
addmeo	0	1	1	1	1	1			rD		rA			///			1	0	1	1	1	0	1	0	1	0	1	0	0	X	addmeo			
addmeo.	0	1	1	1	1	1			rD		rA			///			1	0	1	1	1	0	1	0	1	0	1	0	1	X	addmeo.			
mullwo	0	1	1	1	1	1			rD		rA			rB			1	0	1	1	1	0	1	0	1	0	1	1	0	X	mullwo			
mullwo.	0	1	1	1	1	1			rD		rA			rB			1	0	1	1	1	0	1	0	1	0	1	1	1	X	mullwo.			
dcba <sup>(6)</sup>	0	1	1	1	1	1			///		rA			rB			1	0	1	1	1	1	0	1	1	0	/	X	dcba					
addo	0	1	1	1	1	1			rD		rA			rB			1	1	0	0	0	0	0	1	0	1	0	0	X	addo				
addo.	0	1	1	1	1	1			rD		rA			rB			1	1	0	0	0	0	0	1	0	1	0	1	0	X	addo.			
lhbrx	0	1	1	1	1	1			rD		rA			rB			1	1	0	0	0	1	0	1	1	0	/	X	lhbrx					



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
sraw	0	1	1	1	1	1			rS					rA					rB		1	1	0	0	0	1	1	0	0	0	0	X	sraw	
sraw.	0	1	1	1	1	1			rS					rA					rB		1	1	0	0	0	1	1	0	0	0	1	X	sraw.	
srawi	0	1	1	1	1	1			rS					rA					SH		1	1	0	0	1	1	1	0	0	0	0	X	srawi	
srawi.	0	1	1	1	1	1			rS					rA					SH		1	1	0	0	1	1	1	0	0	0	1	X	srawi.	
sthbrx	0	1	1	1	1	1			rS					rA					rB		1	1	1	0	0	1	0	1	1	0	/	X	sthbrx	
extsh	0	1	1	1	1	1			rS					rA					///		1	1	1	0	0	1	1	0	1	0	0	X	extsh	
extsh.	0	1	1	1	1	1			rS					rA					///		1	1	1	0	0	1	1	0	1	0	1	X	extsh.	
extsb	0	1	1	1	1	1			rS					rA					///		1	1	1	0	1	1	1	0	1	0	0	X	extsb	
extsb.	0	1	1	1	1	1			rS					rA					///		1	1	1	0	1	1	1	0	1	0	1	X	extsb.	
divwuo	0	1	1	1	1	1			rD					rA					rB		1	1	1	1	0	0	1	0	1	1	0	X	divwuo	
divwuo.	0	1	1	1	1	1			rD					rA					rB		1	1	1	1	0	0	1	0	1	1	1	X	divwuo.	
icbi	0	1	1	1	1	1			///					rA					rB		1	1	1	1	0	1	0	1	1	0	/	X	icbi	
divwo	0	1	1	1	1	1			rD					rA					rB		1	1	1	1	1	0	1	0	1	1	0	X	divwo	
divwo.	0	1	1	1	1	1			rD					rA					rB		1	1	1	1	1	0	1	0	1	1	1	X	divwo.	
dcbz	0	1	1	1	1	1			///					rA					rB		1	1	1	1	1	1	0	1	1	0	/	X	dcbz	
lwz	1	0	0	0	0	0			rD					rA																		D	lwz	
lwzu	1	0	0	0	0	1			rD					rA																		D	lwzu	
lbz	1	0	0	0	1	0			rD					rA																		D	lbz	
lbzu	1	0	0	0	1	1			rD					rA																		D	lbzu	
stw	1	0	0	1	0	0			rS					rA																		D	stw	
stwu	1	0	0	1	0	1			rS					rA																		D	stwu	
stb	1	0	0	1	1	0			rS					rA																		D	stb	
stbu	1	0	0	1	1	1			rS					rA																		D	stbu	
lhz	1	0	1	0	0	0			rD					rA																		D	lhz	
lhzu	1	0	1	0	0	1			rD					rA																		D	lhzu	
lha	1	0	1	0	1	0			rD					rA																		D	lha	
lhau	1	0	1	0	1	1			rD					rA																		D	lhau	
sth	1	0	1	1	0	0			rS					rA																		D	sth	



Table 272. Instructions sorted by opcode (binary) (continued)

Mnemonic	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Form	Mnemonic
sth	1	0	1	1	0	1	rS			rA			D															D	sth					
lmw	1	0	1	1	1	0	rD			rA			D															D	lmw					
stmw	1	0	1	1	1	1	rS			rA			D															D	stmw					
fres <sup>(6)</sup>	1	1	1	0	1	1	frD			///			frB			///			1	1	0	0	0	0					A	fres				
fres. <sup>(6)</sup>	1	1	1	0	1	1	frD			///			frB			///			1	1	0	0	0	1					A	fres.				
fsel <sup>(6)</sup>	1	1	1	1	1	1	frD			frA			frB			frC			1	0	1	1	1	0					A	fsel				
fsel. <sup>(6)</sup>	1	1	1	1	1	1	frD			frA			frB			frC			1	0	1	1	1	1					A	fsel.				

1. Supervisor-level instruction.
2. d = UIMM \* 8.
3. d = UIMM \* 2.
4. d = UIMM \* 4.
5. This field is defined as allocated by the Book E architecture, for possible use in an implementation.
6. Optional to the PowerPC classic architecture.
7. Access level is determined by whether the SPR is defined as a user- or supervisor-level SPR.

## A.5 Instruction set legend

Table 273 provides general information on the instruction set (such as architectural level, privilege level, and form).

Table 273. PowerPC instruction set legend

	UIISA	VEA	OEA	Supervisor level	Optional	Form	
addx	√					XO	addx
addcx	√					XO	addcx
addex	√					XO	addex
addi	√					D	addi
addic	√					D	addic
addic.	√					D	addic.
addis	√					D	addis
addmex	√					XO	addmex
addzex	√					XO	addzex
andx	√					X	andx
andcx	√					X	andcx
andi.	√					D	andi.
andis.	√					D	andis.
bx	√					I	bx
bcx	√					B	bcx
bcctrx	√					XL	bcctrx
bclrx	√					XL	bclrx
cmp	√					X	cmp
cmpi	√					D	cmpi
cmpl	√					X	cmpl
cmpli	√					D	cmpli
cntlzwx	√					X	cntlzwx
crand	√					XL	crand
crandc	√					XL	crandc
creqv	√					XL	creqv
crnand	√					XL	crnand
crnor	√					XL	crnor
cror	√					XL	cror
crorc	√					XL	crorc
crxor	√					XL	crxor
dcba		√			√	X	dcba

Table 273. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor level	Optional	Form	
dcbf		√				X	dcbf
dcbi			√	√		X	dcbi
dcbst		√				X	dcbst
dcbt		√				X	dcbt
dcbtst		√				X	dcbtst
dcbz		√				X	dcbz
divwx	√					XO	divwx
divwux	√					XO	divwux
eciwx		√			√	X	eciwx
ecowx		√			√	X	ecowx
eieio		√				X	eieio
eqvx	√					X	eqvx
extsbx	√					X	extsbx
extshx	√					X	extshx
fabsx	√					X	fabsx
faddx	√					A	faddx
faddsx	√					A	faddsx
fcmpo	√					X	fcmpo
fcmpu	√					X	fcmpu
fctiwx	√					X	fctiwx
fctiwzx	√					X	fctiwzx
fdivx	√					A	fdivx
fdivsx	√					A	fdivsx
fmaddx	√					A	fmaddx
fmaddsx	√					A	fmaddsx
fmrx	√					X	fmrx
fmsubx	√					A	fmsubx
fmsubsx	√					A	fmsubsx
fmulx	√					A	fmulx
fmulsx	√					A	fmulsx
fnabsx	√					X	fnabsx
fnegx	√					X	fnegx
fnmaddx	√					A	fnmaddx
fnmaddsx	√					A	fnmaddsx
fnmsubx	√					A	fnmsubx



Table 273. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor level	Optional	Form	
fnmsubsx	√					A	fnmsubsx
fresx	√				√	A	fresx
frspx	√					X	frspx
frsqtex	√				√	A	frsqtex
fselx	√				√	A	fselx
fsqrtx	√				√	A	fsqrtx
fsqrtsx	√				√	A	fsqrtsx
fsubx	√					A	fsubx
fsubsx	√					A	fsubsx
icbi		√				X	icbi
isync		√				XL	isync
lbz	√					D	lbz
lbzu	√					D	lbzu
lbzux	√					X	lbzux
lbzx	√					X	lbzx
lfd	√					D	lfd
lfdx	√					D	lfdx
lfdux	√					X	lfdux
lfdx	√					X	lfdx
lfs	√					D	lfs
lfsu	√					D	lfsu
lfsux	√					X	lfsux
lfsx	√					X	lfsx
lha	√					D	lha
lhau	√					D	lhau
lhaux	√					X	lhaux
lhax	√					X	lhax
lhbrx	√					X	lhbrx
lhz	√					D	lhz
lhzu	√					D	lhzu
lhzux	√					X	lhzux
lhzx	√					X	lhzx
lmw <sup>(1)</sup>	√					D	lmw <sup>(2)</sup>
lswi <sup>(1)</sup>	√					X	lswi <sup>(1)</sup>
lswx <sup>(1)</sup>	√					X	lswx <sup>(1)</sup>

Table 273. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor level	Optional	Form	
lwarx	√					X	lwarx
lwbrx	√					X	lwbrx
lwz	√					D	lwz
lwzu	√					D	lwzu
lwzux	√					X	lwzux
lwzx	√					X	lwzx
mcrf	√					XL	mcrf
mcrfs	√					X	mcrfs
mcrxr	√					X	mcrxr
mfcrr	√					X	mfcrr
mffs	√					X	mffs
mfmsr			√	√		X	mfmsr
mfspr <sup>(3)</sup>	√		√	√		AFX	mfspr <sup>(3)</sup>
mfsr			√	√		X	mfsr
mfsrin			√	√		X	mfsrin
mftb		√				AFX	mftb
mtcrf	√					AFX	mtcrf
mtfsb0x	√					X	mtfsb0x
mtfsb1x	√					X	mtfsb1x
mtfsfx	√					XFL	mtfsfx
mtfsfix	√					X	mtfsfix
mtmsr			√	√		X	mtmsr
mtspr <sup>(3)</sup>	√		√	√		AFX	mtspr <sup>(4)</sup>
mtsr			√	√		X	mtsr
mtsrin			√	√		X	mtsrin
mulhwx	√					XO	mulhwx
mulhwux	√					XO	mulhwux
mulli	√					D	mulli
nandx	√					X	nandx
negx	√					XO	negx
norx	√					X	norx
orx	√					X	orx
orcx	√					X	orcx
ori	√					D	ori
oris	√					D	oris

Table 273. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor level	Optional	Form	
rfi			√	√		XL	rfi
rlwimix	√					M	rlwimix
rlwinmx	√					M	rlwinmx
rlwnmx	√					M	rlwnmx
sc	√		√			SC	sc
slwx	√					X	slwx
srawx	√					X	srawx
srawix	√					X	srawix
srwx	√					X	srwx
stb	√					D	stb
stbu	√					D	stbu
stbux	√					X	stbux
stbx	√					X	stbx
stfd	√					D	stfd
stfdu	√					D	stfdu
stfdux	√					X	stfdux
stfdx	√					X	stfdx
stfiwx	√					X	stfiwx
stfs	√					D	stfs
stfsu	√					D	stfsu
stfsux	√					X	stfsux
stfsx	√					X	stfsx
sth	√					D	sth
sthbrx	√					X	sthbrx
sthu	√					D	sthu
sthux	√					X	sthux
sthx	√					X	sthx
stmw <sup>(1)</sup>	√					D	stmw <sup>(1)</sup>
stswi <sup>(1)</sup>	√					X	stswi <sup>(1)</sup>
stswx <sup>(1)</sup>	√					X	stswx <sup>(1)</sup>
stw	√					D	stw
stwbrx	√					X	stwbrx
stwcx.	√					X	stwcx.
stwu	√					D	stwu
stwux	√					X	stwux

**Table 273. PowerPC instruction set legend (continued)**

	UISA	VEA	OEA	Supervisor level	Optional	Form	
stwx	√					X	stwx
subfx	√					XO	subfx
subfcx	√					XO	subfcx
subfex	√					XO	subfex
subfic	√					D	subfic
subfmex	√					XO	subfmex
subfzex	√					XO	subfzex
sync	√					X	sync
tlbiax			√	√	√	X	tlbiax
tlbiex			√	√	√	X	tlbiex
tlbsync			√	√	√	X	tlbsync
tw	√					X	tw
twi	√					D	twi
xorx	√					X	xorx
xori	√					D	xori
xoris	√					D	xoris

1. Load/Store string or multiple.
2. Load/Store string or multiple.
3. Supervisor and user level instruction.
4. Supervisor and user level instruction.

**Table 274. PowerPC instruction set legend**

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>addx</b>	√					XO	<b>addx</b>
<b>addcx</b>	√					XO	<b>addcx</b>
<b>addex</b>	√					XO	<b>addex</b>
<b>addi</b>	√					D	<b>addi</b>
<b>addic</b>	√					D	<b>addic</b>
<b>addic.</b>	√					D	<b>addic.</b>
<b>addis</b>	√					D	<b>addis</b>
<b>addmex</b>	√					XO	<b>addmex</b>
<b>addzex</b>	√					XO	<b>addzex</b>
<b>andx</b>	√					X	<b>andx</b>
<b>andcx</b>	√					X	<b>andcx</b>
<b>andi.</b>	√					D	<b>andi.</b>



Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>andis.</b>	√					D	<b>andis.</b>
<b>bx</b>	√					I	<b>bx</b>
<b>bcx</b>	√					B	<b>bcx</b>
<b>bcctrx</b>	√					XL	<b>bcctrx</b>
<b>bclrx</b>	√					XL	<b>bclrx</b>
<b>cmp</b>	√					X	<b>cmp</b>
<b>cmpi</b>	√					D	<b>cmpi</b>
<b>cmpl</b>	√					X	<b>cmpl</b>
<b>cmpli</b>	√					D	<b>cmpli</b>
<b>cntlzwx</b>	√					X	<b>cntlzwx</b>
<b>crand</b>	√					XL	<b>crand</b>
<b>crandc</b>	√					XL	<b>crandc</b>
<b>creqv</b>	√					XL	<b>creqv</b>
<b>crnand</b>	√					XL	<b>crnand</b>
<b>crnor</b>	√					XL	<b>crnor</b>
<b>cror</b>	√					XL	<b>cror</b>
<b>crorc</b>	√					XL	<b>crorc</b>
<b>crxor</b>	√					XL	<b>crxor</b>
<b>dcba</b>		√			√	X	<b>dcba</b>
<b>dcbf</b>		√				X	<b>dcbf</b>
<b>dcbi</b>			√	√		X	<b>dcbi</b>
<b>dcbst</b>		√				X	<b>dcbst</b>
<b>dcbt</b>		√				X	<b>dcbt</b>
<b>dcbstst</b>		√				X	<b>dcbstst</b>
<b>dcbz</b>		√				X	<b>dcbz</b>
<b>divwx</b>	√					XO	<b>divwx</b>
<b>divwux</b>	√					XO	<b>divwux</b>
<b>eciwx</b>		√			√	X	<b>eciwx</b>
<b>ecowx</b>		√			√	X	<b>ecowx</b>
<b>eieio</b>		√				X	<b>eieio</b>
<b>eqvx</b>	√					X	<b>eqvx</b>
<b>extsbx</b>	√					X	<b>extsbx</b>
<b>extshx</b>	√					X	<b>extshx</b>
<b>fabsx</b>	√					X	<b>fabsx</b>
<b>faddx</b>	√					A	<b>faddx</b>

Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>faddsx</b>	√					A	<b>faddsx</b>
<b>fcmpo</b>	√					X	<b>fcmpo</b>
<b>fcmpu</b>	√					X	<b>fcmpu</b>
<b>fctiw<sub>x</sub></b>	√					X	<b>fctiw<sub>x</sub></b>
<b>fctiwz<sub>x</sub></b>	√					X	<b>fctiwz<sub>x</sub></b>
<b>fdiv<sub>x</sub></b>	√					A	<b>fdiv<sub>x</sub></b>
<b>fdivs<sub>x</sub></b>	√					A	<b>fdivs<sub>x</sub></b>
<b>fmadd<sub>x</sub></b>	√					A	<b>fmadd<sub>x</sub></b>
<b>fmaddsx</b>	√					A	<b>fmaddsx</b>
<b>fmr<sub>x</sub></b>	√					X	<b>fmr<sub>x</sub></b>
<b>fmsub<sub>x</sub></b>	√					A	<b>fmsub<sub>x</sub></b>
<b>fmsubs<sub>x</sub></b>	√					A	<b>fmsubs<sub>x</sub></b>
<b>fmul<sub>x</sub></b>	√					A	<b>fmul<sub>x</sub></b>
<b>fmuls<sub>x</sub></b>	√					A	<b>fmuls<sub>x</sub></b>
<b>fnabs<sub>x</sub></b>	√					X	<b>fnabs<sub>x</sub></b>
<b>fneg<sub>x</sub></b>	√					X	<b>fneg<sub>x</sub></b>
<b>fnmadd<sub>x</sub></b>	√					A	<b>fnmadd<sub>x</sub></b>
<b>fnmaddsx<sub>x</sub></b>	√					A	<b>fnmaddsx<sub>x</sub></b>
<b>fnmsub<sub>x</sub></b>	√					A	<b>fnmsub<sub>x</sub></b>
<b>fnmsubs<sub>x</sub></b>	√					A	<b>fnmsubs<sub>x</sub></b>
<b>fres<sub>x</sub></b>	√				√	A	<b>fres<sub>x</sub></b>
<b>frsp<sub>x</sub></b>	√					X	<b>frsp<sub>x</sub></b>
<b>frsqrte<sub>x</sub></b>	√				√	A	<b>frsqrte<sub>x</sub></b>
<b>fsel<sub>x</sub></b>	√				√	A	<b>fsel<sub>x</sub></b>
<b>fsqrt<sub>x</sub></b>	√				√	A	<b>fsqrt<sub>x</sub></b>
<b>fsqrts<sub>x</sub></b>	√				√	A	<b>fsqrts<sub>x</sub></b>
<b>fsub<sub>x</sub></b>	√					A	<b>fsub<sub>x</sub></b>
<b>fsubs<sub>x</sub></b>	√					A	<b>fsubs<sub>x</sub></b>
<b>icbi</b>		√				X	<b>icbi</b>
<b>isync</b>		√				XL	<b>isync</b>
<b>lbz</b>	√					D	<b>lbz</b>
<b>lbzu</b>	√					D	<b>lbzu</b>
<b>lbzux</b>	√					X	<b>lbzux</b>

Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
lbzx	√					X	lbzx
lfd	√					D	lfd
lfdu	√					D	lfdu
lfdux	√					X	lfdux
lfdx	√					X	lfdx
lfs	√					D	lfs
lfsu	√					D	lfsu
lfsux	√					X	lfsux
lfsx	√					X	lfsx
lha	√					D	lha
lhau	√					D	lhau
lhaux	√					X	lhaux
lhax	√					X	lhax
lhbrx	√					X	lhbrx
lhz	√					D	lhz
lhzu	√					D	lhzu
lhzux	√					X	lhzux
lhzx	√					X	lhzx
lmw <sup>1</sup>	√					D	lmw <sup>1</sup>
lswi <sup>1</sup>	√					X	lswi <sup>1</sup>
lswx <sup>1</sup>	√					X	lswx <sup>1</sup>
lwarx	√					X	lwarx
lwbrx	√					X	lwbrx
lwz	√					D	lwz
lwzu	√					D	lwzu
lwzux	√					X	lwzux
lwzx	√					X	lwzx
mcrf	√					XL	mcrf
mcrfs	√					X	mcrfs
mcrxr	√					X	mcrxr
mfcrr	√					X	mfcrr
mffs	√					X	mffs
mfmsr			√	√		X	mfmsr
mfmspr <sup>3</sup>	√		√	√		XFX	mfmspr <sup>1</sup>
mfsr			√	√		X	mfsr

Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>mfsrin</b>			√	√		X	<b>mfsrin</b>
<b>mftb</b>		√				AFX	<b>mftb</b>
<b>mterf</b>	√					AFX	<b>mterf</b>
<b>mtfsb0x</b>	√					X	<b>mtfsb0x</b>
<b>mtfsb1x</b>	√					X	<b>mtfsb1x</b>
<b>mtfsfx</b>	√					AXL	<b>mtfsfx</b>
<b>mtfsfix</b>	√					X	<b>mtfsfix</b>
<b>mtmsr</b>			√	√		X	<b>mtmsr</b>
<b>mtspr<sup>1</sup></b>	√		√	√		AFX	<b>mtspr<sup>1</sup></b>
<b>mtsr</b>			√	√		X	<b>mtsr</b>
<b>mtsrin</b>			√	√		X	<b>mtsrin</b>
<b>mulhwx</b>	√					XO	<b>mulhwx</b>
<b>mulhwux</b>	√					XO	<b>mulhwux</b>
<b>mulli</b>	√					D	<b>mulli</b>
<b>mullwx</b>	√					XO	<b>mullwx</b>
<b>nandx</b>	√					X	<b>nandx</b>
<b>negx</b>	√					XO	<b>negx</b>
<b>norx</b>	√					X	<b>norx</b>
<b>orx</b>	√					X	<b>orx</b>
<b>orcx</b>	√					X	<b>orcx</b>
<b>ori</b>	√					D	<b>ori</b>
<b>oris</b>	√					D	<b>oris</b>
<b>rfi</b>			√	√		XL	<b>rfi</b>
<b>rlwimix</b>	√					M	<b>rlwimix</b>
<b>rlwinmx</b>	√					M	<b>rlwinmx</b>
<b>rlwnmx</b>	√					M	<b>rlwnmx</b>
<b>sc</b>	√		√			SC	<b>sc</b>
<b>slwx</b>	√					X	<b>slwx</b>
<b>srawx</b>	√					X	<b>srawx</b>
<b>srawix</b>	√					X	<b>srawix</b>
<b>srwx</b>	√					X	<b>srwx</b>
<b>stb</b>	√					D	<b>stb</b>
<b>stbu</b>	√					D	<b>stbu</b>
<b>stbux</b>	√					X	<b>stbux</b>
<b>stbx</b>	√					X	<b>stbx</b>



Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>stfd</b>	√					D	<b>stfd</b>
<b>stfdu</b>	√					D	<b>stfdu</b>
<b>stfdux</b>	√					X	<b>stfdux</b>
<b>stfdx</b>	√					X	<b>stfdx</b>
<b>stfiwx</b>	√					X	<b>stfiwx</b>
<b>stfs</b>	√					D	<b>stfs</b>
<b>stfsu</b>	√					D	<b>stfsu</b>
<b>stfsux</b>	√					X	<b>stfsux</b>
<b>stfsx</b>	√					X	<b>stfsx</b>
<b>sth</b>	√					D	<b>sth</b>
<b>sthbrx</b>	√					X	<b>sthbrx</b>
<b>sthu</b>	√					D	<b>sthu</b>
<b>sthux</b>	√					X	<b>sthux</b>
<b>sthx</b>	√					X	<b>sthx</b>
<b>stmw</b> <sup>1</sup>	√					D	<b>stmw</b> <sup>1</sup>
<b>stswi</b> <sup>1</sup>	√					X	<b>stswi</b> <sup>1</sup>
<b>stswx</b> <sup>1</sup>	√					X	<b>stswx</b> <sup>1</sup>
<b>stw</b>	√					D	<b>stw</b>
<b>stwbrx</b>	√					X	<b>stwbrx</b>
<b>stwcx.</b>	√					X	<b>stwcx.</b>
<b>stwu</b>	√					D	<b>stwu</b>
<b>stwux</b>	√					X	<b>stwux</b>
<b>stwx</b>	√					X	<b>stwx</b>
<b>subfx</b>	√					XO	<b>subfx</b>
<b>subfcx</b>	√					XO	<b>subfcx</b>
<b>subfex</b>	√					XO	<b>subfex</b>
<b>subfic</b>	√					D	<b>subfic</b>
<b>subfmex</b>	√					XO	<b>subfmex</b>
<b>subfzex</b>	√					XO	<b>subfzex</b>
<b>sync</b>	√					X	<b>sync</b>
<b>tlbiax</b>			√	√	√	X	<b>tlbiax</b>
<b>tlbiex</b>			√	√	√	X	<b>tlbiex</b>
<b>tlbsync</b>			√	√	√	X	<b>tlbsync</b>
<b>tw</b>	√					X	<b>tw</b>
<b>twi</b>	√					D	<b>twi</b>

Table 274. PowerPC instruction set legend (continued)

	UISA	VEA	OEA	Supervisor Level	Optional	Form	
<b>xorx</b>	√					X	<b>xorx</b>
<b>xori</b>	√					D	<b>xori</b>
<b>xoris</b>	√					D	<b>xoris</b>

## Appendix B Simplified mnemonics for PowerPC instructions

This chapter describes simplified mnemonics, which are provided for easier coding of assembly language programs. Simplified mnemonics are defined for the most frequently used forms of branch conditional, compare, trap, rotate and shift, and certain other instructions defined by the PowerPC™ architecture and by implementations of and extensions to the PowerPC architecture.

*B.11: Comprehensive list of simplified mnemonics*, provides an alphabetical listing of simplified mnemonics. Some assemblers may define additional simplified mnemonics not included here. The simplified mnemonics listed here should be supported by all compilers.

### B.1 Overview

Simplified (or extended) mnemonics allow an assembly-language programmer to program using more intuitive mnemonics and symbols than the instructions and syntax defined by the instruction set architecture. For example, to code the conditional call “branch to an absolute target if CR4 specifies a greater than condition, setting the LR without simplified mnemonics, the programmer would write the branch conditional instruction, **bc 12,17,target**. The simplified mnemonic, branch if greater than, **bgt cr4, target**, incorporates the conditions. Not only is it easier to remember the symbols than the numbers when programming, it is also easier to interpret simplified mnemonics when reading existing code.

Although the original PowerPC architecture documents include a set of simplified mnemonics, these are not a formal part of the architecture, but rather a recommendation for assemblers that support the instruction set.

Many simplified mnemonics have been added to those originally included in the architecture documentation. Some assemblers created their own, and others have been added to support extensions to the instruction set (for example, AltiVec instructions and Book E auxiliary processing units (APUs)). Simplified mnemonics have been added for new architecturally defined and new implementation-specific special-purpose registers (SPRs). These simplified mnemonics are described only in a very general way.

### B.2 Subtract simplified mnemonics

This section describes simplified mnemonics for subtract instructions.

#### B.2.1 Subtract immediate

There is no subtract immediate instruction, however, its effect is achieved by negating the immediate operand of an Add Immediate instruction, **addi**. Simplified mnemonics include this negation, making the intent of the computation more clear. These are listed in *Table 275*.

Table 275. Subtract immediate simplified mnemonics

Simplified mnemonic	Standard mnemonic
subi rD,rA,value	addi rD,rA,-value
subis rD,rA,value	addis rD,rA,-value
subic rD,rA,value	addic rD,rA,-value
subic. rD,rA,value	addic. rD,rA,-value

## B.2.2 Subtract

Subtract from instructions subtract the second operand (rA) from the third (rB). The simplified mnemonics in [Table 276](#) use the common order in which the third operand is subtracted from the second.

Table 276. Subtract simplified mnemonics

Simplified mnemonic	Standard mnemonic <sup>(1)</sup>
sub[o][.] rD,rA,rB	subf[o][.] rD,rB,rA
subc[o][.] rD,rA,rB	subfc[o][.] rD,rB,rA

1. rD,rB,rA is not the standard order for the operands. The order of rB and rA is reversed to show the equivalent behavior of the simplified mnemonic.

## B.3 Rotate and shift simplified mnemonics

Rotate and shift instructions provide powerful, general ways to manipulate register contents, but can be difficult to understand. Simplified mnemonics are provided for the following operations:

- Extract—Select a field of  $n$  bits starting at bit position  $b$  in the source register; left or right justify this field in the target register; clear all other bits of the target register.
- Insert—Select a left- or right-justified field of  $n$  bits in the source register; insert this field starting at bit position  $b$  of the target register; leave other bits of the target register unchanged.
- Rotate—Rotate the contents of a register right or left  $n$  bits without masking.
- Shift—Shift the contents of a register right or left  $n$  bits, clearing vacated bits (logical shift).
- Clear—Clear the leftmost or rightmost  $n$  bits of a register.
- Clear left and shift left—Clear the leftmost  $b$  bits of a register, then shift the register left by  $n$  bits. This operation can be used to scale a (known non-negative) array index by the width of an element.

### B.3.1 Operations on words

The simplified mnemonics in [Table 277](#) can be coded with a dot (.) suffix to cause the Rc bit to be set in the underlying instruction.

Table 277. Word rotate and shift simplified mnemonics

Operation	Simplified mnemonic	Equivalent to:
Extract and left justify word immediate	<b>extlwi</b> <i>rA,rS,n,b</i> ( $n > 0$ )	<b>rlwinm</b> <i>rA,rS,b,0,n-1</i>
Extract and right justify word immediate	<b>extrwi</b> <i>rA,rS,n,b</i> ( $n > 0$ )	<b>rlwinm</b> <i>rA,rS,b+n,32-n,31</i>
Insert from left word immediate	<b>inslwi</b> <i>rA,rS,n,b</i> ( $n > 0$ )	<b>rlwimi</b> <i>rA,rS,32-b,b,(b+n)-1</i>
Insert from right word immediate	<b>insrwi</b> <i>rA,rS,n,b</i> ( $n > 0$ )	<b>rlwimi</b> <i>rA,rS,32-(b+n),b,(b+n)-1</i>
Rotate left word immediate	<b>rotlwi</b> <i>rA,rS,n</i>	<b>rlwinm</b> <i>rA,rS,n,0,31</i>
Rotate right word immediate	<b>rotrwi</b> <i>rA,rS,n</i>	<b>rlwinm</b> <i>rA,rS,32-n,0,31</i>
Rotate word left	<b>rotlw</b> <i>rA,rS,rB</i>	<b>rlwnm</b> <i>rA,rS,rB,0,31</i>
Shift left word immediate	<b>slwi</b> <i>rA,rS,n</i> ( $n < 32$ )	<b>rlwinm</b> <i>rA,rS,n,0,31-n</i>
Shift right word immediate	<b>srwi</b> <i>rA,rS,n</i> ( $n < 32$ )	<b>rlwinm</b> <i>rA,rS,32-n,n,31</i>
Clear left word immediate	<b>clrlwi</b> <i>rA,rS,n</i> ( $n < 32$ )	<b>rlwinm</b> <i>rA,rS,0,n,31</i>
Clear right word immediate	<b>clrrwi</b> <i>rA,rS,n</i> ( $n < 32$ )	<b>rlwinm</b> <i>rA,rS,0,0,31-n</i>
Clear left and shift left word immediate	<b>clrlslwi</b> <i>rA,rS,b,n</i> ( $n \leq b \leq 31$ )	<b>rlwinm</b> <i>rA,rS,n,b-n,31-n</i>

Examples using word mnemonics follow:

1. Extract the sign bit (bit 0) of *rS* and place the result right-justified into *rA*.  
**extrwi** *rA,rS,1,0* equivalent to **rlwinm** *rA,rS,1,31,31*
2. Insert the bit extracted in (1) into the sign bit (bit 0) of *rB*.  
**insrwi** *rB,rA,1,0* equivalent to **rlwimi** *rB,rA,31,0,0*
3. Shift the contents of *rA* left 8 bits.  
**slwi** *rA,rA,8* equivalent to **rlwinm** *rA,rA,8,0,23*
4. Clear the high-order 16 bits of *rS* and place the result into *rA*.  
**clrlwi** *rA,rS,16* equivalent to **rlwinm** *rA,rS,0,16,31*

## B.4 Branch instruction simplified mnemonics

Branch conditional instructions can be coded with the operations, a condition to be tested, and a prediction, as part of the mnemonic rather than as numeric BO and BI operands.

[Table 278](#) shows the four general types of branch instructions. Simplified mnemonics are defined only for branch instructions that include BO and BI operands; there is no need to simplify unconditional branch mnemonics.

Table 278. Branch instructions

Instruction name	Mnemonic	Syntax
Branch	<i>b</i> ( <i>ba bl bla</i> )	<i>target_addr</i>
Branch Conditional	<i>bc</i> ( <i>bca bcl bcla</i> )	<i>BO,BI,target_addr</i>

Table 278. Branch instructions

Instruction name	Mnemonic	Syntax
Branch Conditional to Link Register	bclr (bclrl)	BO,BI
Branch Conditional to Count Register	bcctr (bcctrl)	BO,BI

The BO and BI operands correspond to two fields in the instruction opcode, as figure below shows for Branch Conditional (**bc**, **bca**, **bcl**, and **bcla**) instructions.

0	5 6	10 11	15 16	29 30 31
0 0 1 0 0 0	BO	BI	BD	AA LK

The BO operand specifies branch operations that involve decrementing CTR. It is also used to determine whether testing a CR bit causes a branch to occur if the condition is true or false.

The BI operand identifies a CR bit to test (whether a comparison is less than or greater than, for example). The simplified mnemonics avoid the need to memorize the numerical values for BO and BI.

For example, **bc 16,0,target** is a conditional branch that, as a BO value of 16 (0b1\_0000) indicates, decrements CTR, then branches if the decremented CTR is not zero. The operation specified by BO is abbreviated as **d** (for decrement) and **nz** (for not zero), which replace the **c** in the original mnemonic; so the simplified mnemonic for **bc** becomes **bdnz**. The branch does not depend on a condition in the CR, so BI can be eliminated, reducing the expression to **bdnz target**.

In addition to CTR operations, the BO operand provides an optional prediction bit and a true or false indicator can be added. For example, if the previous instruction should branch only on an equal condition in CR0, the instruction becomes **bc 8,2,target**. To incorporate a true condition, the BO value becomes 8 (as shown in [Table 280](#)); the CR0 equal field is indicated by a BI value of 2 (as shown in [Table 281](#)). Incorporating the branch-if-true condition adds a 't' to the simplified mnemonic, **bdnzt**. The equal condition, that is specified by a BI value of 2 (indicating the EQ bit in CR0) is replaced by the **eq** symbol. Using the simplified mnemonic and the **eq** operand, the expression becomes **bdnzt eq,target**.

This example tests CR0[EQ]; however, to test the equal condition in CR5 (CR bit 22), the expression becomes **bc 8,22,target**. The BI operand of 22 indicates CR[22] (CR5[2], or BI field 0b10110), as shown in [Table 281](#). This can be expressed as the simplified mnemonic. **bdnzt 4 \* cr5 + eq,target**.

The notation, **4 \* cr5 + eq** may at first seem awkward, but it eliminates computing the value of the CR bit. It can be seen that  $(4 * 5) + 2 = 22$ . Note that although 32-bit registers in Book E processors are numbered 32–63, only values 0–31 are valid (or possible) for BI operands. As shown in [Table 282](#), a Book E-compliant processor automatically translates the bit values; specifying a BI value of 22 selects bit 55 on a Book E processor, or  $CR5[2] = CR5[EQ]$ .

### B.4.1 Key facts about simplified branch mnemonics


The following key points are helpful in understanding how to use simplified branch mnemonics:

- All simplified branch mnemonics eliminate the BO operand, so if any operand is present in a branch simplified mnemonic, it is the BI operand (or a reduced form of it).
- If the CR is not involved in the branch, the BI operand can be deleted.
- If the CR is involved in the branch, the BI operand can be treated in the following ways:
  - It can be specified as a numeric value, just as it is in the architecturally defined instruction, or it can be indicated with an easier to remember formula,  $4 * crn + [test\ bit\ symbol]$ , where  $n$  indicates the CR field number.
  - The condition of the test bit (eq, lt, gt, and so) can be incorporated into the mnemonic, leaving the need for an operand that defines only the CR field.
    - If the test bit is in CR0, no operand is needed.
    - If the test bit is in CR1–CR7, the BI operand can be replaced with a **crS** operand (that is, **cr1**, **cr2**, **cr3**, and so forth).

### B.4.2 Eliminating the BO operand

The 5-bit BO field, shown below, encodes the following operations in conditional branch instructions:

- Decrement count register (CTR)
  - And test if result is equal to zero
  - And test if result is not equal to zero
- Test condition register (CR)
  - Test condition true
  - Test condition false
- Branch prediction (taken, fall through). If the prediction bit,  $y$ , is needed, it is signified by appending a plus or minus sign as described in [B.4.3: Incorporating the BO branch prediction](#).



BO bits can be interpreted individually as described in [Table 279](#).

**Table 279. BO bit encodings**

BO Bit	Description
0	If set, ignore the CR bit comparison.
1	If set, the CR bit comparison is against true, if not set the CR bit comparison is against false
2	If set, the CTR is not decremented.

Table 279. BO bit encodings (continued)

BO Bit	Description
3	If BO[2] is set, this bit determines whether the CTR comparison is for equal to zero or not equal to zero.
4	The y bit. If set, reverses the static prediction. Use of this bit is optional and independent from the interpretation of other BO bits. Because simplified branch mnemonics eliminate the BO operand, this bit is programmed by adding a plus or minus sign to the simplified mnemonic, as described in <a href="#">B.4.3: Incorporating the BO branch prediction</a> .

Thus, a BO encoding of 10100 (decimal 20) means ignore the CR bit comparison and do not decrement the CTR—in other words, branch unconditionally. Encodings for the BO operand are shown in [Table 280](#). A z bit indicates that the bit is ignored. However, these bits should be cleared, as they may be assigned a meaning in a future version of the architecture.

As shown in [Table 280](#), the 'c' in the standard mnemonic is replaced with the operations otherwise specified in the BO field, (**d** for decrement, **z** for zero, **nz** for non-zero, **t** for true, and **f** for false).

Table 280. BO operand encodings

BO field	Value <sup>(1)</sup> (decimal)	Description	Symbol
0000y	0	Decrement the CTR, then branch if the decremented CTR $\neq$ 0; condition is FALSE.	<b>dnzf</b>
0001y	2	Decrement the CTR, then branch if the decremented CTR = 0; condition is FALSE.	<b>dzf</b>
001zy	4	Branch if the condition is FALSE. <sup>(2)</sup> Note that 'false' and 'four' both start with 'f'.	<b>f</b>
0100y	8	Decrement the CTR, then branch if the decremented CTR $\neq$ 0; condition is TRUE.	<b>dnzt</b>
0101y	10	Decrement the CTR, then branch if the decremented CTR = 0; condition is TRUE.	<b>dzt</b>
011z <sup>(3)</sup> y	12	Branch if the condition is TRUE. <sup>(2)</sup> Note that 'true' and 'twelve' both start with 't'.	<b>t</b>
1z00y <sup>(4)</sup>	16	Decrement the CTR, then branch if the decremented CTR $\neq$ 0.	<b>dnz<sup>(5)</sup></b>
1z01y <sup>(4)</sup>	18	Decrement the CTR, then branch if the decremented CTR = 0.	<b>dz<sup>(5)</sup></b>
1z1zz <sup>(4)</sup>	20	Branch always.	—

- Assumes y = z = 0. [B.4.3: Incorporating the BO branch prediction](#), describes how to use simplified mnemonics to program the y bit for static prediction.
- Instructions for which BO is 12 (branch if condition true) or 4 (branch if condition false) do not depend on the CTR value and can be alternately coded by incorporating the condition specified by the BI field, as described in [B.4.6: Simplified mnemonics that incorporate CR conditions \(eliminates BO and replaces BI with crS\)](#).
- A z bit indicates a bit that is ignored. However, these bits should be cleared, as they may be assigned a meaning in a future version of the architecture.
- Simplified mnemonics for branch instructions that do not test CR bits (BO = 16, 18, and 20) should specify only a target. Otherwise a programming error may occur.
- Notice that these instructions do not use the branch if condition true or false operations. For that reason, simplified mnemonics for these should not specify a BI operand.



### B.4.3 Incorporating the BO branch prediction

As shown in [Table 280](#), the low-order bit (*y* bit) of the BO field provides a hint about whether the branch is likely to be taken (static branch prediction). Assemblers should clear this bit unless otherwise directed. This default action indicates the following:

- A branch conditional with a negative displacement field is predicted to be taken.
- A branch conditional with a non-negative displacement field is predicted not to be taken (fall through).
- A branch conditional to an address in the LR or CTR is predicted not to be taken (fall through).

If the likely outcome (branch or fall through) of a given branch conditional instruction is known, a suffix can be added to the mnemonic that tells the assembler how to set the *y* bit. That is, '+' indicates that the branch is to be taken and '-' indicates that the branch is not to be taken. This suffix can be added to any branch conditional mnemonic, standard or simplified.

For relative and absolute branches (**bc[l][a]**), the setting of the *y* bit depends on whether the displacement field is negative or non-negative. For negative displacement fields, coding the suffix '+' causes the bit to be cleared, and coding the suffix '-' causes the bit to be set. For non-negative displacement fields, coding the suffix '+' causes the bit to be set, and coding the suffix '-' causes the bit to be cleared.

For branches to an address in the LR or CTR (**bclr[l]** or **bcctr[l]**), coding the suffix '+' causes the *y* bit to be set, and coding the suffix '-' causes the bit to be cleared.

Examples of branch prediction follow:

1. Branch if CR0 reflects less than condition, specifying that the branch should be predicted as taken.

`bbl+ target`

2. Same as (1), but target address is in the LR and the branch should be predicted as not taken.

`bblr-`

### B.4.4 The BI operand—CR bit and field representations

With standard branch mnemonics, the BI operand is used when it is necessary to test a CR bit, as shown in the example in [B.4: Branch instruction simplified mnemonics](#).

With simplified mnemonics, the BI operand is handled differently depending on whether the simplified mnemonic incorporates a CR condition to test, as follows:

- Some branch simplified mnemonics incorporate only the BO operand. These simplified mnemonics can use the architecturally defined BI operand to specify the CR bit, as follows:
  - The BI operand can be presented exactly as it is with standard mnemonics—as a decimal number, 0–31.
  - Symbols can be used to replace the decimal operand, as shown in the example in [B.4: Branch instruction simplified mnemonics](#), where `bdnzt 4 * cr5 + eq, target` could be used instead of `bdnzt 22, target`. This is described in [Specifying a CR bit](#)

on page 970.

The simplified mnemonics in *B.4.5: Simplified mnemonics that incorporate the BO operand*, use one of these two methods to specify a CR bit.

- Additional simplified mnemonics are specified that incorporate CR conditions that would otherwise be specified by the BI operand, so the BI operand is replaced by the **crS** operand to specify the CR field, CR0–CR7. See *BI operand instruction encoding on page 970*.

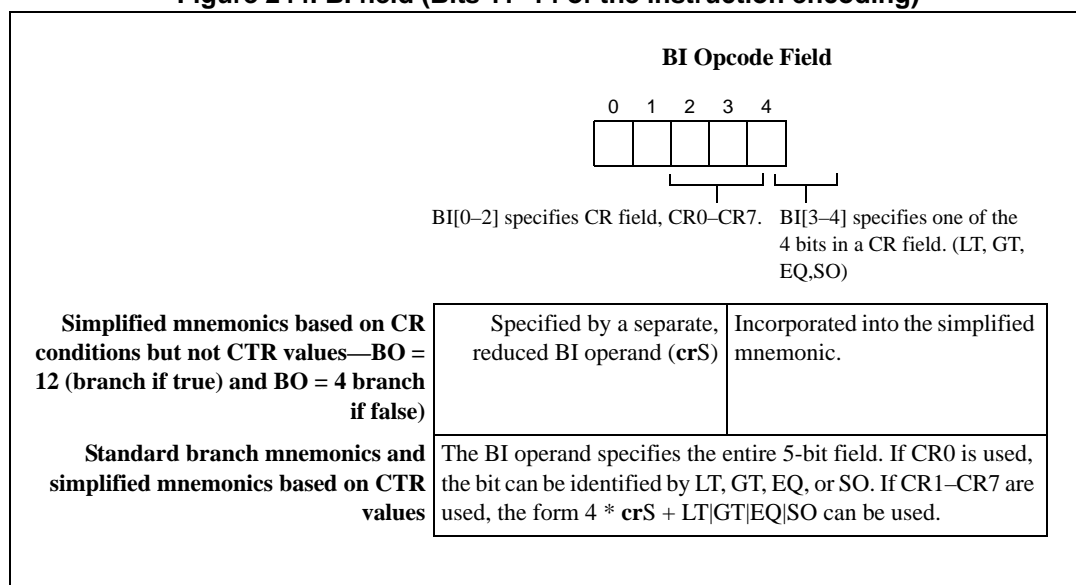
These mnemonics are described in *B.4.6: Simplified mnemonics that incorporate CR conditions (eliminates BO and replaces BI with crS)*.

### BI operand instruction encoding

The entire 5-bit BI field, shown in *Figure 244*, represents the bit number for the CR bit to be tested. For standard branch mnemonics and for branch simplified mnemonics that do not incorporate a CR condition, the BI operand provides all 5 bits.

For simplified branch mnemonics described in *B.4.6*, the BI operand is replaced by a **crS** operand. To understand this, it is useful to view the BI operand as comprised of two parts. As *Figure 244* shows, BI[0–2] indicates the CR field and BI[3–4] represents the condition to test.

**Figure 244. BI field (Bits 11–14 of the instruction encoding)**



Integer record-form instructions update CR0 as described in *Table 281*.

### Specifying a CR bit

Note that the AIM version the PowerPC architecture numbers CR bits 0–31 and Book E numbers them 32–63. However, no adjustment is necessary to the code; in Book E devices, 32 is automatically added to the BI value, as shown in *Table 281* and *Table 282*.

Table 281. CR0 and CR1 fields as updated by integer instructions

CR <sub>n</sub> bit	CR bits		BI		Description
	AIM	Book E	0–2	3–4	
CR0[0]	0	32	000	00	Negative (LT)—Set when the result is negative.
CR0[1]	1	33	000	01	Positive (GT)—Set when the result is positive (and not zero).
CR0[2]	2	34	000	10	Zero (EQ)—Set when the result is zero.
CR0[3]	3	35	000	11	Summary overflow (SO). Copy of XER[SO] at the instruction's completion.

Some simplified mnemonics incorporate only the BO field (as described [B.4.2: Eliminating the BO operand](#)). If one of these simplified mnemonics is used and the CR must be accessed, the BI operand can be specified either as a numeric value or by using the symbols in [Table 282](#).

Compare word instructions (described in [B.5: Compare word simplified mnemonics](#)), move to CR instructions, and others can also modify CR fields, so CR0 and CR1 may hold values that do not adhere to the meanings described in [Table 281](#). CR logical instructions, described in [B.6: Condition register logical simplified mnemonics](#), can update individual CR bits.

Table 282. BI operand settings for CR fields for branch comparisons

CR <sub>n</sub> bit	Bit expression	CR Bits		BI		Description
		AIM (BI operand)	Book E	0–2	3–4	
CR <sub>n</sub> [0]	4 * cr0 + lt (or lt)	0	32	000	00	Less than (LT). For integer compare instructions: rA < SIMM or rB (signed comparison) or rA < UIMM or rB (unsigned comparison).
	4 * cr1 + lt	4	36	001		
	4 * cr2 + lt	8	40	010		
	4 * cr3+ lt	12	44	011		
	4 * cr4 + lt	16	48	100		
	4 * cr5 + lt	20	52	101		
	4 * cr6 + lt	24	56	110		
	4 * cr7 + lt	28	60	111		
CR <sub>n</sub> [1]	4 * cr0 + gt (or gt)	1	33	000	01	Greater than (GT). For integer compare instructions: rA > SIMM or rB (signed comparison) or rA > UIMM or rB (unsigned comparison).
	4 * cr1 + gt	5	37	001		
	4 * cr2 + gt	9	41	010		
	4 * cr3+ gt	13	45	011		
	4 * cr4 + gt	17	49	100		
	4 * cr5 + gt	21	53	101		
	4 * cr6 + gt	25	57	110		
	4 * cr7 + gt	29	61	111		

Table 282. BI operand settings for CR fields for branch comparisons (continued)

CRn bit	Bit expression	CR Bits		BI		Description
		AIM (BI operand)	Book E	0–2	3–4	
CRn[2]	4 * cr0 + eq (or eq)	2	34	000	10	Equal (EQ). For integer compare instructions: rA = SIMM, UIMM, or rB.
	4 * cr1 + eq	6	38	001		
	4 * cr2 + eq	10	42	010		
	4 * cr3 + eq	14	46	011		
	4 * cr4 + eq	18	50	100		
	4 * cr5 + eq	22	54	101		
	4 * cr6 + eq	26	58	110		
	4 * cr7 + eq	30	62	111		
CRn[3]	4 * cr0 + so (or so)	3	35	000	11	Summary overflow (SO). For integer compare instructions, this is a copy of XER[SO] at instruction completion.
	4 * cr1 + so	7	39	001		
	4 * cr2 + so	11	43	010		
	4 * cr3 + so	15	47	011		
	4 * cr4 + so	19	51	100		
	4 * cr5 + so	23	55	101		
	4 * cr6 + so	27	59	110		
	4 * cr7 + so	31	63	111		

To provide simplified mnemonics for every possible combination of BO and BI (that is, including bits that identified the CR field) would require  $2^{10} = 1024$  mnemonics, most of that would be only marginally useful. The abbreviated set in [B.4.5: Simplified mnemonics that incorporate the BO operand](#), covers useful cases. Unusual cases can be coded using a standard branch conditional syntax.

The **crS** operand

The **crS** symbols are shown in [Table 283](#). Note that either the symbol or the operand value can be used in the syntax used with the simplified mnemonic.

Table 283. CR field identification symbols

Symbol	BI[0–2]	CR bits
<b>cr0</b> (default, can be eliminated from syntax)	000	32–35
<b>cr1</b>	001	36–39
<b>cr2</b>	010	40–43
<b>cr3</b>	011	44–47
<b>cr4</b>	100	48–51
<b>cr5</b>	101	52–55
<b>cr6</b>	110	56–59
<b>cr7</b>	111	60–63

To identify a CR bit, an expression in which a CR field symbol is multiplied by 4 and then added to a bit-number-within-CR-field symbol can be used, (for example, **cr0 \* 4 + eq**).

### B.4.5 Simplified mnemonics that incorporate the BO operand

The mnemonics in [Table 284](#) allow common BO operand encodings to be specified as part of the mnemonic, along with the absolute address (AA) and set link register bits (LK). There are no simplified mnemonics for relative and absolute unconditional branches. For these, the basic mnemonics **b**, **ba**, **bl**, and **bla** are used.

**Table 284. Branch simplified mnemonics**

Branch semantics	LR update not enabled				LR update enabled			
	bc	bca	bclr	bcctr	bcl	bcla	bclrl	bcctrl
Branch unconditionally <sup>(1)</sup>	—	—	blr	bctr	—	—	blrl	bctrl
Branch if condition true	bt	bta	btlr	btctr	btl	btla	btlrl	btctrl
Branch if condition false	bf	bfa	bflr	bfctr	bfl	bfla	bflrl	bfctrl
Decrement CTR, branch if CTR ≠ 0 <sup>(1)</sup>	bdnz	bdnza	bdnzlr	—	bdnzl	bdnzla	bdnzlrl	—
Decrement CTR, branch if CTR ≠ 0 and condition true	bdnzt	bdnzta	bdnztlr	—	bdnztl	bdnztla	bdnztlrl	—
Decrement CTR, branch if CTR ≠ 0 and condition false	bdnzf	bdnzfa	bdnzflr	—	bdnzfl	bdnzfla	bdnzflrl	—
Decrement CTR, branch if CTR = 0 <sup>(1)</sup>	bdz	bdza	bdzlr	—	bdzl	bdzla	bdzlrl	—
Decrement CTR, branch if CTR = 0 and condition true	bdzt	bdzta	bdztlr	—	bdztl	bdztla	bdztlrl	—
Decrement CTR, branch if CTR = 0 and condition false	bdzf	bdzfa	bdzflr	—	bdzfl	bdzfla	bdzflrl	—

1. Simplified mnemonics for branch instructions that do not test CR bits should specify only a target. Otherwise a programming error may occur.

[Table 284](#) shows the syntax for basic simplified branch mnemonics

**Table 285. Branch instructions**

Instruction	Standard mnemonic	Syntax	Simplified mnemonic	Syntax
Branch	b (ba bl bla)	target_addr	N/A, syntax does not include BO	
Branch Conditional	bc (bca bcl bcla)	BO, BI, target_addr	<b>bx<sup>(1)</sup></b> (bxa bxl bxla)	BI <sup>(2)</sup> target_addr
Branch Conditional to Link Register	bclr (bclrl)	BO, BI	bxlr (bxlrl)	BI
Branch Conditional to Count Register	bcctr (bcctrl)	BO, BI	bxctr (bxctrl)	BI

- x stands for one of the symbols in [Table 280](#), where applicable.
- BI can be a numeric value or an expression as shown in [Table 283](#).

The simplified mnemonics in [Table 284](#) that test a condition require a corresponding CR bit as the first operand (as examples <sup>(2)</sup>–<sup>(5)</sup> below show). The symbols in [Table 283](#) can be substituted for numeric values.

### Examples that eliminate the BO operand

The simplified mnemonics in [Table 284](#) are used in the following examples:

- Decrement CTR and branch if it is still nonzero (closure of a loop controlled by a count loaded into CTR) (note that no CR bits are tested).  
**bdnz target** equivalent to **bc 16,0,target**  
Because this instruction does not test a CR bit, the simplified mnemonic should specify only a target operand. Specifying a CR (for example, **bdnz 0,target** or **bdnz cr0,target**) may be considered a programming error. Subsequent examples test conditions).
- Same as (1) but branch only if CTR is nonzero and equal condition in CR0.  
**bdnzt eq,target** equivalent to **bc 8,2,target**  
Other equivalents include **bdnzt 2,target** or the unlikely **bdnzt 4\*cr0+eq,target**
- Same as (2), but equal condition is in CR5.  
**bdnzt 4 \* cr5 + eq,target** equivalent to **bc 8,22,target**  
**bdnzt 22,target** would also work
- Branch if bit 59 of CR is false.  
**bf 27,target** equivalent to **bc 4,27,target**  
**bf 4\*cr6+so,target** would also work
- Same as (4), but set the link register. This is a form of conditional call.  
**bfl 27,target** equivalent to **bc 4,27,target**

[Table 286](#) lists simplified mnemonics and syntax for **bc** and **bca** without LR updating.

**Table 286. Simplified mnemonics for bc and bca without LR update**

Branch semantics	bc	Simplified mnemonic	bca	Simplified mnemonic
Branch unconditionally	—	—	—	—
Branch if condition true <sup>(1)</sup>	<b>bc 12,BI,target</b>	<b>bt BI,target</b>	<b>bca 12,BI,target</b>	<b>bta BI,target</b>
Branch if condition false <sup>(1)</sup>	<b>bc 4,BI,target</b>	<b>bf BI,target</b>	<b>bca 4,BI,target</b>	<b>bfa BI,target</b>
Decrement CTR, branch if CTR ≠ 0	<b>bc 16,0,target</b>	<b>bdnz target</b> <sup>(2)</sup>	<b>bca 16,0,target</b>	<b>bdnza target</b> <sup>(2)</sup>
Decrement CTR, branch if CTR ≠ 0 and condition true	<b>bc 8,BI,target</b>	<b>bdnzt BI,target</b>	<b>bca 8,BI,target</b>	<b>bdnzta BI,target</b>
Decrement CTR, branch if CTR ≠ 0 and condition false	<b>bc 0,BI,target</b>	<b>bdnzf BI,target</b>	<b>bca 0,BI,target</b>	<b>bdnzfa BI,target</b>
Decrement CTR, branch if CTR = 0	<b>bc 18,0,target</b>	<b>bdz target</b> <sup>(2)</sup>	<b>bca 18,0,target</b>	<b>bdza target</b> <sup>(2)</sup>
Decrement CTR, branch if CTR = 0 and condition true	<b>bc 10,BI,target</b>	<b>bdzt BI,target</b>	<b>bca 10,BI,target</b>	<b>bdzta BI,target</b>
Decrement CTR, branch if CTR = 0 and condition false	<b>bc 2,BI,target</b>	<b>bdzf BI,target</b>	<b>bca 2,BI,target</b>	<b>bdzfa BI,target</b>

- Instructions for which BO is either 12 (branch if condition true) or 4 (branch if condition false) do not depend on the CTR value and can be alternately coded by incorporating the condition specified by the BI field, as described in [B.4.6: Simplified mnemonics that incorporate CR conditions \(eliminates BO and replaces BI with crS\)](#).
- Simplified mnemonics for branch instructions that do not test CR bits should specify only a target. Otherwise a programming error may occur.

[Table 287](#) lists simplified mnemonics and syntax for **bclr** and **bcctr** without LR updating.

Table 287. Simplified mnemonics for bclr and bcctr without LR update

Branch Semantics	bclr	Simplified mnemonic	bcctr	Simplified mnemonic
Branch unconditionally	<b>bclr 20,0</b>	<b>blr</b> <sup>(1)</sup>	<b>bcctr 20,0</b>	<b>bctr</b> <sup>(1)</sup>
Branch if condition true <sup>(2)</sup>	<b>bclr 12,BI</b>	<b>btlr BI</b>	<b>bcctr 12,BI</b>	<b>btctr BI</b>
Branch if condition false <sup>(2)</sup>	<b>bclr 4,BI</b>	<b>bftr BI</b>	<b>bcctr 4,BI</b>	<b>bfctr BI</b>
Decrement CTR, branch if CTR ≠ 0	<b>bclr 16,BI</b>	<b>bdnzlr BI</b>	—	—
Decrement CTR, branch if CTR ≠ 0 and condition true	<b>bclr 8,BI</b>	<b>bdnztlr BI</b>	—	—
Decrement CTR, branch if CTR ≠ 0 and condition false	<b>bclr 0,BI</b>	<b>bdnzflr BI</b>	—	—
Decrement CTR, branch if CTR = 0	<b>bclr 18,0</b>	<b>bdzlr</b> <sup>(1)</sup>	—	—
Decrement CTR, branch if CTR = 0 and condition true	<b>bclr 8,BI</b>	<b>bdnztlr BI</b>	—	—
Decrement CTR, branch if CTR = 0 and condition false	<b>bclr 2,BI</b>	<b>bdzflr BI</b>	—	—

1. Simplified mnemonics for branch instructions that do not test a CR bit should not specify one; a programming error may occur.
2. Instructions for which B0 is 12 (branch if condition true) or 4 (branch if condition false) do not depend on a CTR value and can be alternately coded by incorporating the condition specified by the BI field. See [B.4.6: Simplified mnemonics that incorporate CR conditions \(eliminates B0 and replaces BI with crS\)](#).

[Table 288](#) provides simplified mnemonics and syntax for **bcl** and **bcla**.

Table 288. Simplified mnemonics for bcl and bcla with LR update

Branch semantics	bcl	Simplified mnemonic	bcla	Simplified mnemonic
Branch unconditionally	—	—	—	—
Branch if condition true <sup>(1)</sup>	<b>bcl 12,BI,target</b>	<b>btl BI,target</b>	<b>bcla 12,BI,target</b>	<b>btla BI,target</b>
Branch if condition false <sup>(1)</sup>	<b>bcl 4,BI,target</b>	<b>bfl BI,target</b>	<b>bcla 4,BI,target</b>	<b>bfla BI,target</b>
Decrement CTR, branch if CTR ≠ 0	<b>bcl 16,0,target</b>	<b>bdnzl target</b> <sup>(2)</sup>	<b>bcla 16,0,target</b>	<b>bdnzla target</b> <sup>(2)</sup>
Decrement CTR, branch if CTR ≠ 0 and condition true	<b>bcl 8,0,target</b>	<b>bdnztl BI,target</b>	<b>bcla 8,BI,target</b>	<b>bdnzla BI,target</b>
Decrement CTR, branch if CTR ≠ 0 and condition false	<b>bcl 0,BI,target</b>	<b>bdnzfl BI,target</b>	<b>bcla 0,BI,target</b>	<b>bdnzfla BI,target</b>
Decrement CTR, branch if CTR = 0	<b>bcl 18,BI,target</b>	<b>bdzl target</b> <sup>(2)</sup>	<b>bcla 18,BI,target</b>	<b>bdzla target</b> <sup>(2)</sup>
Decrement CTR, branch if CTR = 0 and condition true	<b>bcl 10,BI,target</b>	<b>bdztl BI,target</b>	<b>bcla 10,BI,target</b>	<b>bdzla BI,target</b>
Decrement CTR, branch if CTR = 0 and condition false	<b>bcl 2,BI,target</b>	<b>bdzfl BI,target</b>	<b>bcla 2,BI,target</b>	<b>bdzfla BI,target</b>

1. Instructions for which B0 is either 12 (branch if condition true) or 4 (branch if condition false) do not depend on the CTR value and can be alternately coded by incorporating the condition specified by the BI field. See [B.4.6: Simplified mnemonics that incorporate CR conditions \(eliminates B0 and replaces BI with crS\)](#).
2. Simplified mnemonics for branch instructions that do not test CR bits should specify only a target. A programming error may occur.

Table 289 provides simplified mnemonics and syntax for **bclrl** and **bcctrl** with LR updating.

**Table 289. Simplified mnemonics for bclrl and bcctrl with LR update**

Branch semantics	bclrl	Simplified mnemonic	bcctrl	simplified mnemonic
Branch unconditionally	<b>bclrl 20,0</b>	<b>blr<sup>(1)</sup></b>	<b>bcctrl 20,0</b>	<b>bctrl<sup>(1)</sup></b>
Branch if condition true	<b>bclrl 12,BI</b>	<b>btlrl BI</b>	<b>bcctrl 12,BI</b>	<b>btctrl BI</b>
Branch if condition false	<b>bclrl 4,BI</b>	<b>bfirl BI</b>	<b>bcctrl 4,BI</b>	<b>bfctrl BI</b>
Decrement CTR, branch if CTR $\neq$ 0	<b>bclrl 16,0</b>	<b>bdnzlrl<sup>(1)</sup></b>	—	—
Decrement CTR, branch if CTR $\neq$ 0, condition true	<b>bclrl 8,BI</b>	<b>bdnztlrl BI</b>	—	—
Decrement CTR, branch if CTR $\neq$ 0, condition false	<b>bclrl 0,BI</b>	<b>bdnzflrl BI</b>	—	—
Decrement CTR, branch if CTR = 0	<b>bclrl 18,0</b>	<b>bdzrlrl<sup>(1)</sup></b>	—	—
Decrement CTR, branch if CTR = 0, condition true	<b>bclrl 10, BI</b>	<b>bdztlrl BI</b>	—	—
Decrement CTR, branch if CTR = 0, condition false	<b>bclrl 2,BI</b>	<b>bdzflrl BI</b>	—	—

1. Simplified mnemonics for branch instructions that do not test a CR bit should not specify one. A programming error may occur.

#### B.4.6 Simplified mnemonics that incorporate CR conditions (eliminates BO and replaces BI with crS)

The mnemonics in Table 292 are variations of the branch-if-condition-true (BO = 12) and branch-if-condition-false (BO = 4) encodings. Because these instructions do not depend on the CTR, the true/false conditions specified by BO can be combined with the CR test bit specified by BI to create a different set of simplified mnemonics that eliminates the BO operand and the portion of the BI operand (BI[3–4]) that specifies one of the four possible test bits. However, the simplified mnemonic cannot specify in which of the eight CR fields the test bit falls, so the BI operand is replaced by a **crS** operand.

The standard codes shown in Table 290 are used for the most common combinations of branch conditions. Note that for ease of programming, these codes include synonyms; for example, less than or equal (**le**) and not greater than (**ng**) achieve the same result.

Note: A CR field symbol, **cr0–cr7**, is used as the first operand after the simplified mnemonic. If **CR0** is used, no **crS** is necessary.

**Table 290. Standard coding for branch conditions**

Code	Description	Equivalent	Bit tested
<b>lt</b>	Less than	—	LT
<b>le</b>	Less than or equal (equivalent to <b>ng</b> )	<b>ng</b>	GT
<b>eq</b>	Equal	—	EQ
<b>ge</b>	Greater than or equal (equivalent to <b>nl</b> )	<b>nl</b>	LT
<b>gt</b>	Greater than	—	GT
<b>nl</b>	Not less than (equivalent to <b>ge</b> )	<b>ge</b>	LT
<b>ne</b>	Not equal	—	EQ



**Table 290. Standard coding for branch conditions (continued)**

Code	Description	Equivalent	Bit tested
<b>ng</b>	Not greater than (equivalent to <b>le</b> )	<b>le</b>	GT
<b>so</b>	Summary overflow	—	SO
<b>ns</b>	Not summary overflow	—	SO

*Table 291* shows the syntax for simplified branch mnemonics that incorporate CR conditions. Here, **crS** replaces a BI operand to specify only a CR field (because the specific CR bit within the field is now part of the simplified mnemonic. Note that the default is CR0; if no **crS** is specified, CR0 is used.

**Table 291. Branch instructions and simplified mnemonics that incorporate CR conditions**

Instruction	Standard mnemonic	Syntax	Simplified mnemonic	Syntax
Branch	b (ba bl bla)	target_addr	—	
Branch Conditional	bc (bca bcl bcla)	BO,BI,target_addr	<b>bx</b> <sup>(1)</sup> ( <b>bx</b> a <b>bx</b> l <b>bx</b> la)	<b>crS</b> <sup>(2)</sup> ,target_addr
Branch Conditional to Link Register	bclr (bclrl)	BO,BI	<b>bx</b> lr ( <b>bx</b> lrl)	<b>crS</b>
Branch Conditional to Count Register	bcctr (bcctrl)	BO,BI	<b>bx</b> ctr ( <b>bx</b> ctrl)	<b>crS</b>

1. x stands for one of the symbols in *Table 290*, where applicable.
2. BI can be a numeric value or an expression as shown in *Table 283*.

*Table 292* shows the simplified branch mnemonics incorporating conditions.

**Table 292. Simplified mnemonics with comparison conditions**

Branch semantics	LR update not enabled				LR update enabled			
	bc	bca	bclr	bcctr	bcl	bcla	bclrl	bcctrl
Branch if less than	<b>blt</b>	<b>blta</b>	<b>bltlr</b>	<b>bltctr</b>	<b>bltl</b>	<b>bltla</b>	<b>bltlrl</b>	<b>bltctrl</b>
Branch if less than or equal	<b>ble</b>	<b>blea</b>	<b>blelr</b>	<b>blectr</b>	<b>blel</b>	<b>blela</b>	<b>blelrl</b>	<b>blectrl</b>
Branch if equal	<b>beq</b>	<b>beqa</b>	<b>beqlr</b>	<b>beqctr</b>	<b>beql</b>	<b>beqla</b>	<b>beqlrl</b>	<b>beqctrl</b>
Branch if greater than or equal	<b>bge</b>	<b>bgea</b>	<b>bgehr</b>	<b>bgectr</b>	<b>bgel</b>	<b>bgeha</b>	<b>bgehr</b>	<b>bgectrl</b>
Branch if greater than	<b>bgt</b>	<b>bgta</b>	<b>bgtr</b>	<b>bgctr</b>	<b>bgtl</b>	<b>bgtha</b>	<b>bgtr</b>	<b>bgctrl</b>
Branch if not less than	<b>bnl</b>	<b>bnla</b>	<b>bnlhr</b>	<b>bnlctr</b>	<b>bnll</b>	<b>bnlla</b>	<b>bnlhr</b>	<b>bnlctrl</b>
Branch if not equal	<b>bne</b>	<b>bnea</b>	<b>bnelr</b>	<b>bnectr</b>	<b>bnel</b>	<b>bnela</b>	<b>bnelrl</b>	<b>bnectrl</b>
Branch if not greater than	<b>bng</b>	<b>bnga</b>	<b>bnglr</b>	<b>bngctr</b>	<b>bngl</b>	<b>bngla</b>	<b>bnglrl</b>	<b>bngctrl</b>
Branch if summary overflow	<b>bso</b>	<b>bsoa</b>	<b>bsolr</b>	<b>bsoctr</b>	<b>bsol</b>	<b>bsola</b>	<b>bsolrl</b>	<b>bsoctrl</b>
Branch if not summary overflow	<b>bns</b>	<b>bnsa</b>	<b>bnslr</b>	<b>bnsctr</b>	<b>bns</b>	<b>bnsa</b>	<b>bnsrl</b>	<b>bnsctrl</b>
Branch if unordered	<b>bun</b>	<b>buna</b>	<b>bunlr</b>	<b>bunctr</b>	<b>bun</b>	<b>buna</b>	<b>bunrl</b>	<b>bunctrl</b>
Branch if not unordered	<b>bnu</b>	<b>bnu</b>	<b>bnulr</b>	<b>bnuctr</b>	<b>bnu</b>	<b>bnu</b>	<b>bnulrl</b>	<b>bnuctrl</b>



Instructions using the mnemonics in [Table 292](#) indicate the condition bit, but not the CR field. If no field is specified, CR0 is used. The CR field symbols defined in [Table 283](#) (**cr0**–**cr7**) are used for this operand, as shown in examples <sup>(2)</sup>–<sup>(4)</sup> below.

### Branch simplified mnemonics that incorporate CR conditions: examples

The following examples use the simplified mnemonics shown in [Table 292](#):

- Branch if CR0 reflects not-equal condition.  
**bne target** equivalent to **bc 4,2,target**
- Same as (1) but condition is in CR3.  
**bne cr3,target** equivalent to **bc 4,14,target**
- Branch to an absolute target if CR4 specifies greater than condition, setting the LR. This is a form of conditional call.  
**bgtla cr4,target** equivalent to **bcla 12,17,target**
- Same as (3), but target address is in the CTR.  
**bgtctrl cr4** equivalent to **bctrl 12,17**

### Branch simplified mnemonics that incorporate CR conditions: listings

[Table 293](#) shows simplified branch mnemonics and syntax for **bc** and **bca** without LR updating.

**Table 293. Simplified mnemonics for bc and bca without comparison conditions or LR Update**

Branch Semantics	bc	Simplified mnemonic	bca	Simplified mnemonic
Branch if less than	<b>bc 12,BI<sup>(1)</sup>,target</b>	<b>blt crS target</b>	<b>bca 12,BI<sup>(1)</sup>,target</b>	<b>blta crS target</b>
Branch if less than or equal	<b>bc 4,BI<sup>(2)</sup>,target</b>	<b>ble crS target</b>	<b>bca 4,BI<sup>(2)</sup>,target</b>	<b>blea crS target</b>
Branch if not greater than		<b>bng crS target</b>		<b>bnga crS target</b>
Branch if equal	<b>bc 12,BI<sup>(3)</sup>,target</b>	<b>beq crS target</b>	<b>bca 12,BI<sup>(3)</sup>,target</b>	<b>beqa crS target</b>
Branch if greater than or equal	<b>bc 4,BI<sup>(1)</sup>,target</b>	<b>bge crS target</b>	<b>bca 4,BI<sup>(1)</sup>,target</b>	<b>bgea crS target</b>
Branch if not less than		<b>bni crS target</b>		<b>bnla crS target</b>
Branch if greater than	<b>bc 12,BI<sup>(2)</sup>,target</b>	<b>bgt crS target</b>	<b>bca 12,BI<sup>(2)</sup>,target</b>	<b>bgta crS target</b>
Branch if not equal	<b>bc 4,BI<sup>(3)</sup>,target</b>	<b>bne crS target</b>	<b>bca 4,BI<sup>(3)</sup>,target</b>	<b>bnea crS target</b>
Branch if summary overflow	<b>bc 12,BI<sup>(4)</sup>,target</b>	<b>bso crS target</b>	<b>bca 12,BI<sup>(4)</sup>,target</b>	<b>bsoa crS target</b>
Branch if unordered		<b>bun crS target</b>		<b>buna crS target</b>
Branch if not summary overflow	<b>bc 4,BI<sup>(4)</sup>,target</b>	<b>bns crS target</b>	<b>bca 4,BI<sup>(4)</sup>,target</b>	<b>bnsa crS target</b>
Branch if not unordered		<b>bnu crS target</b>		<b>bnua crS target</b>

- The value in the BI operand selects CR $n$ [0], the LT bit.
- The value in the BI operand selects CR $n$ [1], the GT bit.
- The value in the BI operand selects CR $n$ [2], the EQ bit.

4. The value in the BI operand selects CRn[3], the SO bit.

Table 294 shows simplified branch mnemonics and syntax for **bclr** and **bcctr** without LR updating.

**Table 294. Simplified mnemonics for bclr and bcctr without comparison conditions or LR update**

Branch semantics	bclr	Simplified mnemonic	bcctr	Simplified mnemonic
Branch if less than	<b>bclr 12,BI<sup>(1)</sup>,target</b>	<b>bltlr crS target</b>	<b>bcctr 12,BI<sup>(1)</sup>,target</b>	<b>bltctr crS target</b>
Branch if less than or equal	<b>bclr 4,BI<sup>(2)</sup>,target</b>	<b>blelr crS target</b>	<b>bcctr 4,BI<sup>(2)</sup>,target</b>	<b>blectr crS target</b>
Branch if not greater than		<b>bnglr crS target</b>		<b>bngctr crS target</b>
Branch if equal	<b>bclr 12,BI<sup>(3)</sup>,target</b>	<b>beqlr crS target</b>	<b>bcctr 12,BI<sup>(3)</sup>,target</b>	<b>beqctr crS target</b>
Branch if greater than or equal	<b>bclr 4,BI<sup>(1)</sup>,target</b>	<b>bgelr crS target</b>	<b>bcctr 4,BI<sup>(1)</sup>,target</b>	<b>bgectr crS target</b>
Branch if not less than		<b>bnllr crS target</b>		<b>bnlctr crS target</b>
Branch if greater than	<b>bclr 12,BI<sup>(2)</sup>,target</b>	<b>bgtlr crS target</b>	<b>bcctr 12,BI<sup>(2)</sup>,target</b>	<b>bgtctr crS target</b>
Branch if not equal	<b>bclr 4,BI<sup>(3)</sup>,target</b>	<b>bnelr crS target</b>	<b>bcctr 4,BI<sup>(3)</sup>,target</b>	<b>bnectr crS target</b>
Branch if summary overflow	<b>bclr 12,BI<sup>(4)</sup>,target</b>	<b>bsolr crS target</b>	<b>bcctr 12,BI<sup>(4)</sup>,target</b>	<b>bsoctr crS target</b>
Branch if not summary overflow	<b>bclr 4,BI<sup>(4)</sup>,target</b>	<b>bnslr crS target</b>	<b>bcctr 4,BI<sup>(4)</sup>,target</b>	<b>bnsctr crS target</b>

1. The value in the BI operand selects CRn[0], the LT bit.
2. The value in the BI operand selects CRn[1], the GT bit.
3. The value in the BI operand selects CRn[2], the EQ bit.
4. The value in the BI operand selects CRn[3], the SO bit.

Table 295 shows simplified branch mnemonics and syntax for **bcl** and **bcla**.

**Table 295. Simplified mnemonics for bcl and bcla with comparison conditions, LR update**

Branch semantics	bcl	Simplified mnemonic	bcla	Simplified mnemonic
Branch if less than	<b>bcl 12,BI<sup>(1)</sup>,target</b>	<b>bltl crS target</b>	<b>bcla 12,BI<sup>(1)</sup>,target</b>	<b>bltla crS target</b>
Branch if less than or equal	<b>bcl 4,BI<sup>(2)</sup>,target</b>	<b>blel crS target</b>	<b>bcla 4,BI<sup>(2)</sup>,target</b>	<b>blela crS target</b>
Branch if not greater than		<b>bngl crS target</b>		<b>bngla crS target</b>
Branch if equal	<b>bcl 12,BI<sup>(3)</sup>,target</b>	<b>beql crS target</b>	<b>bcla 12,BI<sup>(3)</sup>,target</b>	<b>beqla crS target</b>
Branch if greater than or equal	<b>bcl 4,BI<sup>(1)</sup>,target</b>	<b>bgel crS target</b>	<b>bcla 4,BI<sup>(1)</sup>,target</b>	<b>bgela crS target</b>
Branch if not less than		<b>bnll crS target</b>		<b>bnlla crS target</b>
Branch if greater than	<b>bcl 12,BI<sup>(2)</sup>,target</b>	<b>bgtl crS target</b>	<b>bcla 12,BI<sup>(2)</sup>,target</b>	<b>bgtla crS target</b>
Branch if not equal	<b>bcl 4,BI<sup>(3)</sup>,target</b>	<b>bnel crS target</b>	<b>bcla 4,BI<sup>(3)</sup>,target</b>	<b>bnela crS target</b>

**Table 295. Simplified mnemonics for bcl and bcla with comparison conditions, LR update**

Branch semantics	bcl	Simplified mnemonic	bcla	Simplified mnemonic
Branch if summary overflow	<b>bcl 12,BI<sup>(4)</sup>,target</b>	<b>bsol crS target</b>	<b>bcla 12,BI<sup>(4)</sup>,target</b>	<b>bsola crS target</b>
Branch if not summary overflow	<b>bcl 4,BI<sup>(4)</sup>,target</b>	<b>bnsl crS target</b>	<b>bcla 4,BI<sup>(4)</sup>,target</b>	<b>bnsla crS target</b>

1. The value in the BI operand selects CR $n$ [0], the LT bit.
2. The value in the BI operand selects CR $n$ [1], the GT bit.
3. The value in the BI operand selects CR $n$ [2], the EQ bit.
4. The value in the BI operand selects CR $n$ [3], the SO bit.

[Table 296](#) shows the simplified branch mnemonics and syntax for **bclrl** and **bcctrl** with LR updating.

**Table 296. Simplified mnemonics for bclrl and bcctrl with comparison conditions, LR update**

Branch semantics	bclrl	Simplified mnemonic	bcctrl	Simplified mnemonic
Branch if less than	<b>bclrl 12,BI<sup>(1)</sup>,target</b>	<b>blttrl crS target</b>	<b>bcctrl 12,BI<sup>(1)</sup>,target</b>	<b>bltctrl crS target</b>
Branch if less than or equal	<b>bclrl 4,BI<sup>(2)</sup>,target</b>	<b>blelrl crS target</b>	<b>bcctrl 4,BI<sup>(2)</sup>,target</b>	<b>blectrl crS target</b>
Branch if not greater than		<b>bnglrl crS target</b>		<b>bngctrl crS target</b>
Branch if equal	<b>bclrl 12,BI<sup>(3)</sup>,target</b>	<b>beqlrl crS target</b>	<b>bcctrl 12,BI<sup>(3)</sup>,target</b>	<b>beqctrl crS target</b>
Branch if greater than or equal	<b>bclrl 4,BI<sup>(1)</sup>,target</b>	<b>bgeirl crS target</b>	<b>bcctrl 4,BI<sup>(1)</sup>,target</b>	<b>bgectrl crS target</b>
Branch if not less than		<b>bnlrl crS target</b>		<b>bnlctrl crS target</b>
Branch if greater than	<b>bclrl 12,BI<sup>(2)</sup>,target</b>	<b>bgtlrl crS target</b>	<b>bcctrl 12,BI<sup>(2)</sup>,target</b>	<b>bgtctrl crS target</b>
Branch if not equal	<b>bclrl 4,BI<sup>(3)</sup>,target</b>	<b>bnelrl crS target</b>	<b>bcctrl 4,BI<sup>(3)</sup>,target</b>	<b>bnctrl crS target</b>
Branch if summary overflow	<b>bclrl 12,BI<sup>(4)</sup>,target</b>	<b>bsolrl crS target</b>	<b>bcctrl 12,BI<sup>(4)</sup>,target</b>	<b>bsoctrl crS target</b>
Branch if not summary overflow	<b>bclrl 4,BI<sup>(4)</sup>,target</b>	<b>bnsrlrl crS target</b>	<b>bcctrl 4,BI<sup>(4)</sup>,target</b>	<b>bnsctrl crS target</b>

1. The value in the BI operand selects CR $n$ [0], the LT bit.
2. The value in the BI operand selects CR $n$ [1], the GT bit.
3. The value in the BI operand selects CR $n$ [2], the EQ bit.
4. The value in the BI operand selects CR $n$ [3], the SO bit.

## B.5 Compare word simplified mnemonics

In compare word instructions, the L operand indicates a word (L = 0) or a double-word (L = 1). Simplified mnemonics in [Table 297](#) eliminate the L operand for word comparisons.

**Table 297. Word compare simplified mnemonics**

Operation	Simplified mnemonic	Equivalent to:
Compare Word Immediate	<b>cmpwi crD,rA,SIMM</b>	<b>cmpi crD,0,rA,SIMM</b>
Compare Word	<b>cmpw crD,rA,rB</b>	<b>cmp crD,0,rA,rB</b>

Table 297. Word compare simplified mnemonics (continued)

Operation	Simplified mnemonic	Equivalent to:
Compare Logical Word Immediate	cmplwi crD,rA,UIMM	cmpli crD,0,rA,UIMM
Compare Logical Word	cmplw crD,rA,rB	cmpl crD,0,rA,rB

As with branch mnemonics, the **crD** field of a compare instruction can be omitted if CR0 is used, as shown in examples <sup>(1)</sup> and <sup>(3)</sup> below. Otherwise, the target CR field must be specified as the first operand. The following examples use word compare mnemonics:

1. Compare **rA** with immediate value 100 as signed 32-bit integers and place result in CR0.  
**cmpwi rA,100** equivalent to **cmpi 0,0,rA,100**
2. Same as (1), but place results in CR4.  
**cmpwi cr4,rA,100** equivalent to **cmpi 4,0,rA,100**
3. Compare **rA** and **rB** as unsigned 32-bit integers and place result in CR0.  
**cmplw rA,rB** equivalent to **cmpl 0,0,rA,rB**

## B.6 Condition register logical simplified mnemonics

The CR logical instructions, shown in [Table 298](#), can be used to set, clear, copy, or invert a given CR bit. Simplified mnemonics allow these operations to be coded easily. Note that the symbols defined in [Table 282](#) can be used to identify the CR bit.

Table 298. Condition register logical simplified mnemonics

Operation	Simplified mnemonic	Equivalent to
Condition register set	<b>crset bx</b>	<b>creqv bx,bx,bx</b>
Condition register clear	<b>crclr bx</b>	<b>crxor bx,bx,bx</b>
Condition register move	<b>crmove bx,by</b>	<b>cror bx,by,by</b>
Condition register not	<b>crnot bx,by</b>	<b>crnor bx,by,by</b>

Examples using the CR logical mnemonics follow:

1. Set CR[57].  
**crset 25** equivalent to **creqv 25,25,25**
2. Clear CR0[SO].  
**crclr so** equivalent to **crxor 3,3,3**
3. Same as (2), but clear CR3[SO].  
**crclr 4 \* cr3 + so** equivalent to **crxor 15,15,15**
4. Invert the CR0[EQ].  
**crnot eq,eq** equivalent to **crnor 2,2,2**
5. Same as (4), but CR4[EQ] is inverted and the result is placed into CR5[EQ].  
**crnot 4 \* cr5 + eq, 4 \* cr4 + eq** equivalent to **crnor 22,18,18**

## B.7 Trap instructions simplified mnemonics

The codes in [Table 299](#) are for the most common combinations of trap conditions.

**Table 299. Standard codes for trap instructions**

Code	Description	TO encoding	<	>	=	<U <sup>(1)</sup>	>U <sup>(2)</sup>
lt	Less than	16	1	0	0	0	0
le	Less than or equal	20	1	0	1	0	0
eq	Equal	4	0	0	1	0	0
ge	Greater than or equal	12	0	1	1	0	0
gt	Greater than	8	0	1	0	0	0
nl	Not less than	12	0	1	1	0	0
ne	Not equal	24	1	1	0	0	0
ng	Not greater than	20	1	0	1	0	0
llt	Logically less than	2	0	0	0	1	0
lle	Logically less than or equal	6	0	0	1	1	0
lge	Logically greater than or equal	5	0	0	1	0	1
lgt	Logically greater than	1	0	0	0	0	1
lnl	Logically not less than	5	0	0	1	0	1
lng	Logically not greater than	6	0	0	1	1	0
—	Unconditional	31	1	1	1	1	1

1. The symbol '<U' indicates an unsigned less-than evaluation is performed.
2. The symbol '>U' indicates an unsigned greater-than evaluation is performed.

The mnemonics in [Table 300](#) are variations of trap instructions, with the most useful TO values represented in the mnemonic rather than specified as a numeric operand.

**Table 300. Trap simplified mnemonics**

Trap semantics	32-Bit Comparison	
	twi Immediate	tw Register
Trap unconditionally	—	trap
Trap if less than	twlti	twlt
Trap if less than or equal	twlei	twle
Trap if equal	tweqi	tweq
Trap if greater than or equal	twgei	twge
Trap if greater than	twgti	twgt
Trap if not less than	twnli	twnl
Trap if not equal	twnei	twne
Trap if not greater than	twngi	twng

**Table 300. Trap simplified mnemonics (continued)**

Trap semantics	32-Bit Comparison	
	twi Immediate	tw Register
Trap if logically less than	twl $lti$	twl $lt$
Trap if logically less than or equal	twl $lei$	twl $le$
Trap if logically greater than or equal	twl $gei$	twl $ge$
Trap if logically greater than	twl $gti$	twl $gt$
Trap if logically not less than	twl $nli$	twl $nl$
Trap if logically not greater than	twl $ngi$	twl $ng$

The following examples use the trap mnemonics shown in [Table 300](#):

- Trap if rA is not zero.  
twnei rA,0 equivalent to twi 24,rA,0
- Trap if rA is not equal to rB.  
twne rA, rB equivalent to tw 24,rA,rB
- Trap if rA is logically greater than 0x7FF.  
twlgti rA, 0x7FF equivalent to twi 1,rA, 0x7FF
- Trap unconditionally.  
trapequivalent to tw 31,0,0

Trap instructions evaluate a trap condition as follows: The contents of rA are compared with either the sign-extended SIMM field or the contents of rB, depending on the trap instruction.

The comparison results in five conditions that are ANDed with operand TO. If the result is not 0, the trap exception handler is invoked. See [Table 301](#) for these conditions.

**Table 301. TO operand bit encoding**

TO bit	ANDed with condition
0	Less than, using signed comparison
1	Greater than, using signed comparison
2	Equal
3	Less than, using unsigned comparison
4	Greater than, using unsigned comparison

## B.8 Simplified mnemonics for accessing SPRs

The **mtspr** and **mf spr** instructions specify a special-purpose register (SPR) as a numeric operand. Simplified mnemonics are provided that represent the SPR in the mnemonic rather than requiring it to be coded as a numeric operand. The pattern for **mtspr** and **mf spr** simplified mnemonics is straightforward: replace the **-spr** portion of the mnemonic with the abbreviation for the spr (for example XER, SRR0, or LR), eliminate the SPRN operand, leaving the source or destination GPR operand, rS or rD.

Following are examples using the SPR simplified mnemonics:

1. Copy the contents of rS to the XER.  
**mtxer rS** equivalent to **tomtspr 1,rS**
2. Copy the contents of the LR to rD.  
**mflr rD** equivalent to **tomfspr rD,8**
3. Copy the contents of rS to the CTR.  
**mtctr rS** equivalent to **tomtspr 9,rS**

The examples above show simplified mnemonics for accessing SPRs defined by the AIM version of the PowerPC architecture; however, the same formula is used for Book E, EIS, and implementation-specific SPRs, as shown in the following examples:

1. Copy the contents of rS to CSRR0.  
**mtcsrr0 rS** equivalent to **mtspr 58,rS**
2. Copy the contents of IVOR0 to rD.  
**mfivor0 rD** equivalent to **tomfspr rD,400**
3. Copy the contents of rS to the MAS1.  
**mtmas1 rS** equivalent to **tomtspr 625,rS**

There is an additional simplified mnemonic formula for accessing SPRGs, although not all of these more complicated simplified mnemonics are supported by all assemblers. These are shown in [Table 302](#) along with the equivalent simplified mnemonic using the formula described above.

**Table 302. Additional simplified mnemonics for accessing SPRGs**

SPR	Move to SPR		Move from SPR	
	Simplified mnemonic	Equivalent to	Simplified mnemonic	Equivalent to
SPRGs	<b>mtsprg n, rS</b>	<b>mtspr 272 + n,rS</b>	<b>mfsprg rD, n</b>	<b>mfspr rD,272 + n</b>
	<b>mtsprgn, rS</b>		<b>mfsprgn rD</b>	

## B.9 Recommended simplified mnemonics

This section describes commonly-used operations (such as no-op, load immediate, load address, move register, and complement register).

### B.9.1 No-op (nop)

Many instructions can be coded so that, effectively, no operation is performed. A mnemonic is provided for the preferred form of no-op. If an implementation performs any type of run-time optimization related to no-ops, the preferred form is the following:

**nop** equivalent to **ori 0,0,0**

### B.9.2 Load immediate (li)

The **addi** and **addis** instructions can be used to load an immediate value into a register. Additional mnemonics are provided to convey the idea that no addition is being performed but that data is being moved from the immediate operand of the instruction to a register.



1. Load a 16-bit signed immediate value into **rD**.  
**li rD,value** equivalent to **addi rD,0,value**
2. Load a 16-bit signed immediate value, shifted left by 16 bits, into **rD**.  
**lis rD,value** equivalent to **addis rD,0,value**

### B.9.3 Load address (**la**)

This mnemonic permits computing the value of a base-displacement operand, using the **addi** instruction that normally requires a separate register and immediate operands.

**la rD,d(rA)** equivalent to **addi rD,rA,d**

The **la** mnemonic is useful for obtaining the address of a variable specified by name, allowing the assembler to supply the base register number and compute the displacement. If the variable *v* is located at offset *dv* bytes from the address in *rV*, and the assembler has been told to use *rV* as a base for references to the data structure containing *v*, the following line causes the address of *v* to be loaded into **rD**:

**la rD,v** equivalent to **addi rD,rV,dv**

### B.9.4 Move register (**mr**)

Several instructions can be coded to copy the contents of one register to another. A simplified mnemonic is provided that signifies that no computation is being performed, but merely that data is being moved from one register to another.

The following instruction copies the contents of **rS** into **rA**. This mnemonic can be coded with a dot (.) suffix to cause the Rc bit to be set in the underlying instruction.

**mr rA,rS** equivalent to **or rA,rS,rS**

### B.9.5 Complement register (**not**)

Several instructions can be coded in a way that they complement the contents of one register and place the result into another register. Simplified mnemonics allows this operation to be coded easily.

The following instruction complements the contents of **rS** and places the result into **rA**. This mnemonic can be coded with a dot (.) suffix to cause the Rc bit to be set in the underlying instruction.

**not rA,rS** equivalent to **nor rA,rS,rS**

### B.9.6 Move to condition register (**mctr**)

This mnemonic permits copying GPR contents to the CR, using the same syntax as the **mfcrr** instruction.

**mctr rS** equivalent to **mtrcrf 0xFF,rS**

## B.10 EIS-specific simplified mnemonics

This section describes simplified mnemonics for instructions defined by auxiliary processing units (APUs) defined as part of the Motorola Book E implementation standards.

### B.10.1 Integer select (isel)

The following mnemonics simplify the most common variants of the **isel** instruction that access CR0:

Integer Select Less Than <b>isellt</b> rD,rA,rB	equivalent to	<b>isel</b> rD,rA,rB,0
Integer Select Greater Than <b>iselgt</b> rD,rA,rB	equivalent to	<b>isel</b> rD,rA,rB,1
Integer Select Equal <b>iseleq</b> rD,rA,rB	equivalent to	<b>isel</b> rD,rA,rB,2

### B.10.2 SPE mnemonics

The following mnemonic handles moving of the full 64-bit SPE GPR:

Vector Move <b>evmr</b> rD,rA	equivalent to	<b>evor</b> rD,rA,rA
----------------------------------	---------------	----------------------

The following mnemonic performs a complement register:

Vector Not <b>evnot</b> rD,rA	equivalent to	<b>evnor</b> rD,rA,rA
----------------------------------	---------------	-----------------------

## B.11 Comprehensive list of simplified mnemonics

[Table 303](#) lists simplified mnemonics. Note that compiler designers may implement additional simplified mnemonics not listed here.

**Table 303. Simplified mnemonics**

Simplified mnemonic	Mnemonic	Instruction
<b>bctr</b> <sup>(1)</sup>	<b>bcctr 20,0</b>	Branch unconditionally ( <b>bcctr</b> without LR update)
<b>bctrl</b> <sup>(1)</sup>	<b>bcctrl 20,0</b>	Branch unconditionally ( <b>bcctrl</b> with LR Update)
<b>bdnz</b> target <sup>(1)</sup>	<b>bc 16,0,target</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bc</b> without LR update)
<b>bdnza</b> target <sup>(1)</sup>	<b>bca 16,0,target</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bca</b> without LR update)
<b>bdnzf</b> BI,target	<b>bc 0,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bc</b> without LR update)
<b>bdnzfa</b> BI,target	<b>bca 0,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bca</b> without LR update)
<b>bdnzfl</b> BI,target	<b>bcl 0,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bcl</b> with LR update)
<b>bdnzfla</b> BI,target	<b>bcla 0,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bcla</b> with LR update)
<b>bdnzflr</b> BI	<b>bclr 0,BI</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bclr</b> without LR update)
<b>bdnzflrl</b> BI	<b>bclrl 0,BI</b>	Decrement CTR, branch if CTR ≠ 0 and condition false ( <b>bclrl</b> with LR Update)
<b>bdnzl</b> target <sup>(1)</sup>	<b>bcl 16,0,target</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bcl</b> with LR update)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bdnzla</b> target <sup>(1)</sup>	<b>bcla 16,0,target</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bcla</b> with LR update)
<b>bdnzlr BI</b>	<b>bclr 16,BI</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bclr</b> without LR update)
<b>bdnzlrl</b> <sup>(1)</sup>	<b>bclrl 16,0</b>	Decrement CTR, branch if CTR ≠ 0 ( <b>bclrl</b> with LR Update)
<b>bdnzt BI,target</b>	<b>bc 8,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bc</b> without LR update)
<b>bdnzta BI,target</b>	<b>bca 8,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bca</b> without LR update)
<b>bdnztl BI,target</b>	<b>bcl 8,0,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bcl</b> with LR update)
<b>bdnztla BI,target</b>	<b>bcla 8,BI,target</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bcla</b> with LR update)
<b>bdnztlr BI</b>	<b>bclr 8,BI</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bclr</b> without LR update)
<b>bdnztlr BI</b>	<b>bclr 8,BI</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bclr</b> without LR update)
<b>bdnztlrl BI</b>	<b>bclrl 8,BI</b>	Decrement CTR, branch if CTR ≠ 0 and condition true ( <b>bclrl</b> with LR Update)
<b>bdz</b> target <sup>(1)</sup>	<b>bc 18,0,target</b>	Decrement CTR, branch if CTR = 0 ( <b>bc</b> without LR update)
<b>bdza</b> target <sup>(1)</sup>	<b>bca 18,0,target</b>	Decrement CTR, branch if CTR = 0 ( <b>bca</b> without LR update)
<b>bdzf BI,target</b>	<b>bc 2,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bc</b> without LR update)
<b>bdzfa BI,target</b>	<b>bca 2,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bca</b> without LR update)
<b>bdzfl BI,target</b>	<b>bcl 2,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bcl</b> with LR update)
<b>bdzfla BI,target</b>	<b>bcla 2,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bcla</b> with LR update)
<b>bdzflr BI</b>	<b>bclr 2,BI</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bclr</b> without LR update)
<b>bdzflrl BI</b>	<b>bclrl 2,BI</b>	Decrement CTR, branch if CTR = 0 and condition false ( <b>bclrl</b> with LR Update)
<b>bdzl</b> target <sup>(1)</sup>	<b>bcl 18,BI,target</b>	Decrement CTR, branch if CTR = 0 ( <b>bcl</b> with LR update)
<b>bdzla</b> target <sup>(1)</sup>	<b>bcla 18,BI,target</b>	Decrement CTR, branch if CTR = 0 ( <b>bcla</b> with LR update)
<b>bdzlr</b> <sup>(1)</sup>	<b>bclr 18,0</b>	Decrement CTR, branch if CTR = 0 ( <b>bclr</b> without LR update)
<b>bdzlrl</b> <sup>(1)</sup>	<b>bclrl 18,0</b>	Decrement CTR, branch if CTR = 0 ( <b>bclrl</b> with LR Update)
<b>bdzt BI,target</b>	<b>bc 10,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bc</b> without LR update)
<b>bdzta BI,target</b>	<b>bca 10,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bca</b> without LR update)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bdztl</b> BI,target	<b>bcl 10,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bcl</b> with LR update)
<b>bdztl</b> a BI,target	<b>bcla 10,BI,target</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bcla</b> with LR update)
<b>bdztlr</b> l BI	<b>bclrl 10, BI</b>	Decrement CTR, branch if CTR = 0 and condition true ( <b>bclrl</b> with LR Update)
<b>beq</b> crS target	<b>bc 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bc</b> without comparison conditions or LR updating)
<b>beq</b> a crS target	<b>bca 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bca</b> without comparison conditions or LR updating)
<b>beqctr</b> crS target	<b>bcctr 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bcctr</b> without comparison conditions and LR updating)
<b>beqctrl</b> crS target	<b>bcctrl 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bcctrl</b> with comparison conditions and LR update)
<b>beql</b> crS target	<b>bcl 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bcl</b> with comparison conditions and LR updating)
<b>beql</b> a crS target	<b>bcla 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bcla</b> with comparison conditions and LR updating)
<b>beqlr</b> crS target	<b>bclr 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bclr</b> without comparison conditions and LR updating)
<b>beqlrl</b> crS target	<b>bclrl 12,BI<sup>(2)</sup>,target</b>	Branch if equal ( <b>bclrl</b> with comparison conditions and LR update)
<b>bf</b> BI,target	<b>bc 4,BI,target</b>	Branch if condition false <sup>(3)</sup> ( <b>bc</b> without LR update)
<b>bfa</b> BI,target	<b>bca 4,BI,target</b>	Branch if condition false <sup>(3)</sup> ( <b>bca</b> without LR update)
<b>bfctr</b> BI	<b>bcctr 4,BI</b>	Branch if condition false <sup>(3)</sup> ( <b>bcctr</b> without LR update)
<b>bfctrl</b> BI	<b>bcctrl 4,BI</b>	Branch if condition false <sup>(3)</sup> ( <b>bcctrl</b> with LR Update)
<b>bfl</b> BI,target	<b>bcl 4,BI,target</b>	Branch if condition false <sup>(3)</sup> ( <b>bcl</b> with LR update)
<b>bfl</b> a BI,target	<b>bcla 4,BI,target</b>	Branch if condition false <sup>(3)</sup> ( <b>bcla</b> with LR update)
<b>bflr</b> BI	<b>bclr 4,BI</b>	Branch if condition false <sup>(3)</sup> ( <b>bclr</b> without LR update)
<b>bflrl</b> BI	<b>bclrl 4,BI</b>	Branch if condition false <sup>(3)</sup> ( <b>bclrl</b> with LR Update)
<b>bge</b> crS target	<b>bc 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bc</b> without comparison conditions or LR updating)
<b>bge</b> a crS target	<b>bca 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bca</b> without comparison conditions or LR updating)
<b>bgectr</b> crS target	<b>bcctr 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bcctr</b> without comparison conditions and LR updating)
<b>bgectrl</b> crS target	<b>bcctrl 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bge</b> l crS target	<b>bcl 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bcl</b> with comparison conditions and LR updating)
<b>bge</b> la crS target	<b>bcla 4,BI<sup>(4)</sup>,target</b>	Branch if greater than or equal ( <b>bcla</b> with comparison conditions and LR updating)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bgelr</b> crS target	<b>bclr</b> 4,BI <sup>(4)</sup> ,target	Branch if greater than or equal ( <b>bclr</b> without comparison conditions and LR updating)
<b>bgelrl</b> crS target	<b>bclrl</b> 4,BI <sup>(4)</sup> ,target	Branch if greater than or equal ( <b>bclrl</b> with comparison conditions and LR update)
<b>bgt</b> crS target	<b>bc</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bc</b> without comparison conditions or LR updating)
<b>bgta</b> crS target	<b>bca</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bca</b> without comparison conditions or LR updating)
<b>bgtctr</b> crS target	<b>bcctr</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bcctr</b> without comparison conditions and LR updating)
<b>bgtctrl</b> crS target	<b>bcctrl</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bgtl</b> crS target	<b>bcl</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bcl</b> with comparison conditions and LR updating)
<b>bgta</b> crS target	<b>bcla</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bcla</b> with comparison conditions and LR updating)
<b>bgtlr</b> crS target	<b>bclr</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bclr</b> without comparison conditions and LR updating)
<b>bgtlrl</b> crS target	<b>bclrl</b> 12,BI <sup>(5)</sup> ,target	Branch if greater than ( <b>bclrl</b> with comparison conditions and LR update)
<b>ble</b> crS target	<b>bc</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bc</b> without comparison conditions or LR updating)
<b>blea</b> crS target	<b>bca</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bca</b> without comparison conditions or LR updating)
<b>blectr</b> crS target	<b>bcctr</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bcctr</b> without comparison conditions and LR updating)
<b>blectrl</b> crS target	<b>bcctrl</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bcctrl</b> with comparison conditions and LR update)
<b>blel</b> crS target	<b>bcl</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bcl</b> with comparison conditions and LR updating)
<b>blela</b> crS target	<b>bcla</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bcla</b> with comparison conditions and LR updating)
<b>blelr</b> crS target	<b>bclr</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bclr</b> without comparison conditions and LR updating)
<b>blelrl</b> crS target	<b>bclrl</b> 4,BI <sup>(5)</sup> ,target	Branch if less than or equal ( <b>bclrl</b> with comparison conditions and LR update)
<b>blr</b> <sup>(1)</sup>	<b>bclr</b> 20,0	Branch unconditionally ( <b>bclr</b> without LR update)
<b>blr</b> <sup>(1)</sup>	<b>bclrl</b> 20,0	Branch unconditionally ( <b>bclrl</b> with LR Update)
<b>blt</b> crS target	<b>bc</b> 12,BI,target	Branch if less than ( <b>bc</b> without comparison conditions or LR updating)
<b>blta</b> crS target	<b>bca</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bca</b> without comparison conditions or LR updating)
<b>bltctr</b> crS target	<b>bcctr</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bcctr</b> without comparison conditions and LR updating)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bltctrl</b> crS target	<b>bcctrl</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bltl</b> crS target	<b>bcl</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bcl</b> with comparison conditions and LR updating)
<b>bltla</b> crS target	<b>bcla</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bcla</b> with comparison conditions and LR updating)
<b>bltlr</b> crS target	<b>bclr</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bclr</b> without comparison conditions and LR updating)
<b>bltrl</b> crS target	<b>bclrl</b> 12,BI <sup>(4)</sup> ,target	Branch if less than ( <b>bclrl</b> with comparison conditions and LR update)
<b>bne</b> crS target	<b>bc</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bc</b> without comparison conditions or LR updating)
<b>bnea</b> crS target	<b>bca</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bca</b> without comparison conditions or LR updating)
<b>bnctr</b> crS target	<b>bcctr</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bcctr</b> without comparison conditions and LR updating)
<b>bnctrl</b> crS target	<b>bcctrl</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bnel</b> crS target	<b>bcl</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bcl</b> with comparison conditions and LR updating)
<b>bnela</b> crS target	<b>bcla</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bcla</b> with comparison conditions and LR updating)
<b>bnelr</b> crS target	<b>bclr</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bclr</b> without comparison conditions and LR updating)
<b>bnelrl</b> crS target	<b>bclrl</b> 4,BI <sup>(3)</sup> ,target	Branch if not equal ( <b>bclrl</b> with comparison conditions and LR update)
<b>bng</b> crS target	<b>bc</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bc</b> without comparison conditions or LR updating)
<b>bnga</b> crS target	<b>bca</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bca</b> without comparison conditions or LR updating)
<b>bngctr</b> crS target	<b>bcctr</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bcctr</b> without comparison conditions and LR updating)
<b>bngctrl</b> crS target	<b>bcctrl</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bngl</b> crS target	<b>bcl</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bcl</b> with comparison conditions and LR updating)
<b>bngla</b> crS target	<b>bcla</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bcla</b> with comparison conditions and LR updating)
<b>bnglr</b> crS target	<b>bclr</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bclr</b> without comparison conditions and LR updating)
<b>bnglrl</b> crS target	<b>bclrl</b> 4,BI <sup>(5)</sup> ,target	Branch if not greater than ( <b>bclrl</b> with comparison conditions and LR update)
<b>bnl</b> crS target	<b>bc</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bc</b> without comparison conditions or LR updating)
<b>bnla</b> crS target	<b>bca</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bca</b> without comparison conditions or LR updating)
<b>bnlctr</b> crS target	<b>bcctr</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bcctr</b> without comparison conditions and LR updating)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bnlctrl</b> crS target	<b>bcctrl</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bnll</b> crS target	<b>bcl</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bcl</b> with comparison conditions and LR updating)
<b>bnlla</b> crS target	<b>bcla</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bcla</b> with comparison conditions and LR updating)
<b>bnllr</b> crS target	<b>bclr</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bclr</b> without comparison conditions and LR updating)
<b>bnllrl</b> crS target	<b>bclrl</b> 4,BI <sup>(4)</sup> ,target	Branch if not less than ( <b>bclrl</b> with comparison conditions and LR update)
<b>bns</b> crS target	<b>bc</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bc</b> without comparison conditions or LR updating)
<b>bnsa</b> crS target	<b>bca</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bca</b> without comparison conditions or LR updating)
<b>bnsctr</b> crS target	<b>bcctr</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bcctr</b> without comparison conditions and LR updating)
<b>bnsctrl</b> crS target	<b>bcctrl</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bcctrl</b> with comparison conditions and LR update)
<b>bnsll</b> crS target	<b>bcl</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bcl</b> with comparison conditions and LR updating)
<b>bnslla</b> crS target	<b>bcla</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bcla</b> with comparison conditions and LR updating)
<b>bnsllr</b> crS target	<b>bclr</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bclr</b> without comparison conditions and LR updating)
<b>bnsllrl</b> crS target	<b>bclrl</b> 4,BI <sup>(6)</sup> ,target	Branch if not summary overflow ( <b>bclrl</b> with comparison conditions and LR update)
<b>bso</b> crS target	<b>bc</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bc</b> without comparison conditions or LR updating)
<b>soa</b> crS target	<b>bca</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bca</b> without comparison conditions or LR updating)
<b>soctr</b> crS target	<b>bcctr</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bcctr</b> without comparison conditions and LR updating)
<b>soctrl</b> crS target	<b>bcctrl</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bcctrl</b> with comparison conditions and LR update)
<b>sol</b> crS target	<b>bcl</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bcl</b> with comparison conditions and LR updating)
<b>sola</b> crS target	<b>bcla</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bcla</b> with comparison conditions and LR updating)
<b>solr</b> crS target	<b>bclr</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bclr</b> without comparison conditions and LR updating)
<b>solrl</b> crS target	<b>bclrl</b> 12,BI <sup>(6)</sup> ,target	Branch if summary overflow ( <b>bclrl</b> with comparison conditions and LR update)
<b>bt</b> BI,target	<b>bc</b> 12,BI,target	Branch if condition true <sup>(3)</sup> ( <b>bc</b> without LR update)

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>bta</b> BI,target	<b>bca</b> 12,BI,target	Branch if condition true <sup>(3)</sup> ( <b>bca</b> without LR update)
<b>btctr</b> BI	<b>bcctr</b> 12,BI	Branch if condition true <sup>(3)</sup> ( <b>bcctr</b> without LR update)
<b>btctrl</b> BI	<b>bcctrl</b> 12,BI	Branch if condition true <sup>(3)</sup> ( <b>bcctrl</b> with LR Update)
<b>btl</b> BI,target	<b>bcl</b> 12,BI,target	Branch if condition true <sup>(3)</sup> ( <b>bcl</b> with LR update)
<b>btla</b> BI,target	<b>bcla</b> 12,BI,target	Branch if condition true <sup>(3)</sup> ( <b>bcla</b> with LR update)
<b>btlr</b> BI	<b>bclr</b> 12,BI	Branch if condition true <sup>(3)</sup> ( <b>bclr</b> without LR update)
<b>btlrl</b> BI	<b>bclrl</b> 12,BI	Branch if condition true <sup>(3)</sup> ( <b>bclrl</b> with LR Update)
<b>clrlslwi</b> rA,rS,b,n ( $n \leq b \leq 31$ )	<b>rlwinm</b> rA,rS,n,b – n,31 – n	Clear left and shift left word immediate
<b>clrlwi</b> rA,rS,n ( $n < 32$ )	<b>rlwinm</b> rA,rS,0,n,31	Clear left word immediate
<b>clrrwi</b> rA,rS,n ( $n < 32$ )	<b>rlwinm</b> rA,rS,0,0,31 – n	Clear right word immediate
<b>cmplw</b> crD,rA,rB	<b>cmpl</b> crD,0,rA,rB	Compare logical word
<b>cmplwi</b> crD,rA,UIMM	<b>cmpli</b> crD,0,rA,UIMM	Compare logical word immediate
<b>cmpw</b> crD,rA,rB	<b>cmp</b> crD,0,rA,rB	Compare word
<b>cmpwi</b> crD,rA,SIMM	<b>cmpi</b> crD,0,rA,SIMM	Compare word immediate
<b>crclr</b> bx	<b>crxor</b> bx,bx,bx	Condition register clear
<b>crmove</b> bx,by	<b>cror</b> bx,by,by	Condition register move
<b>crnot</b> bx,by	<b>crnor</b> bx,by,by	Condition register not
<b>crset</b> bx	<b>creqv</b> bx,bx,bx	Condition register set
<b>evmr</b> rD,rA	<b>evor</b> rD,rA,rA	Vector Move Register
<b>evnot</b> rD,rA	<b>evnor</b> rD,rA,rA	Vector Complement Register
<b>evsubiw</b> rD,rB,UIMM	<b>evsubifw</b> rD,UIMM,rB	Vector subtract word immediate
<b>evsubw</b> rD,rB,rA	<b>evsubfw</b> rD,rA,rB	Vector subtract word
<b>extlwi</b> rA,rS,n,b ( $n > 0$ )	<b>rlwinm</b> rA,rS,b,0,n – 1	Extract and left justify word immediate
<b>extrwi</b> rA,rS,n,b ( $n > 0$ )	<b>rlwinm</b> rA,rS,b + n, 32 – n,31	Extract and right justify word immediate
<b>inslwi</b> rA,rS,n,b ( $n > 0$ )	<b>rlwimi</b> rA,rS,32 – b,b,(b + n) – 1	Insert from left word immediate
<b>insrwi</b> rA,rS,n,b ( $n > 0$ )	<b>rlwimi</b> rA,rS,32 – (b + n),b,(b + n) – 1	Insert from right word immediate
<b>iseleq</b> rD,rA,rB	<b>isel</b> rD,rA,rB,2	Integer Select Equal



Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>iselgt</b> rD,rA,rB	<b>isel</b> rD,rA,rB,1	Integer Select Greater Than
<b>isellt</b> rD,rA,rB	<b>isel</b> rD,rA,rB,0	Integer Select Less Than
<b>la</b> rD,d(rA)	<b>addi</b> rD,rA,d	Load address
<b>li</b> rD,value	<b>addi</b> rD,0,value	Load immediate
<b>lis</b> rD,value	<b>addis</b> rD,0,value	Load immediate signed
<b>mf spr</b> rD	<b>mf spr</b> rD,SPRN	Move from SPR (see <a href="#">B.8: Simplified mnemonics for accessing SPRs.</a> )
<b>mr</b> rA,rS	<b>or</b> rA,rS,rS	Move register
<b>mtcr</b> rS	<b>mtcrf</b> 0xFF,rS	Move to Condition Register
<b>mt spr</b> rS	<b>mf spr</b> SPRN,rS	Move to SPR (see <a href="#">B.8: Simplified mnemonics for accessing SPRs.</a> )
<b>nop</b>	<b>ori</b> 0,0,0	No-op
<b>not</b> rA,rS	<b>nor</b> rA,rS,rS	NOT
<b>not</b> rA,rS	<b>nor</b> rA,rS,rS	Complement register
<b>rotlw</b> rA,rS,rB	<b>rlwnm</b> rA,rS,rB,0,31	Rotate left word
<b>rotlwi</b> rA,rS,n	<b>rlwinm</b> rA,rS,n,0,31	Rotate left word immediate
<b>rotrwi</b> rA,rS,n	<b>rlwinm</b> rA,rS,32 – n,0,31	Rotate right word immediate
<b>slwi</b> rA,rS,n (n < 32)	<b>rlwinm</b> rA,rS,n,0,31 – n	Shift left word immediate
<b>srwi</b> rA,rS,n (n < 32)	<b>rlwinm</b> rA,rS,32 – n,n,31	Shift right word immediate
<b>sub</b> rD,rA,rB	<b>subf</b> rD,rB,rA	Subtract from
<b>subc</b> rD,rA,rB	<b>subfc</b> rD,rB,rA	Subtract from carrying
<b>subi</b> rD,rA,value	<b>addi</b> rD,rA,–value	Subtract immediate
<b>subic</b> rD,rA,value	<b>addic</b> rD,rA,– value	Subtract immediate carrying
<b>subic.</b> rD,rA,value	<b>addic.</b> rD,rA,– value	Subtract immediate carrying
<b>subis</b> rD,rA,value	<b>addis</b> rD,rA,– value	Subtract immediate signed
<b>tweq</b> rA,SIMM	<b>tw</b> 4,rA,SIMM	Trap if equal
<b>tweqi</b> rA,SIMM	<b>twi</b> 4,rA,SIMM	Trap immediate if equal
<b>twge</b> rA,SIMM	<b>tw</b> 12,rA,SIMM	Trap if greater than or equal
<b>twgei</b> rA,SIMM	<b>twi</b> 12,rA,SIMM	Trap immediate if greater than or equal
<b>twgt</b> rA,SIMM	<b>tw</b> 8,rA,SIMM	Trap if greater than
<b>twgti</b> rA,SIMM	<b>twi</b> 8,rA,SIMM	Trap immediate if greater than

Table 303. Simplified mnemonics (continued)

Simplified mnemonic	Mnemonic	Instruction
<b>twle</b> rA,SIMM	<b>tw 20</b> ,rA,SIMM	Trap if less than or equal
<b>twlei</b> rA,SIMM	<b>twi 20</b> ,rA,SIMM	Trap immediate if less than or equal
<b>twlge</b> rA,SIMM	<b>tw 12</b> ,rA,SIMM	Trap if logically greater than or equal
<b>twlgei</b> rA,SIMM	<b>twi 12</b> ,rA,SIMM	Trap immediate if logically greater than or equal
<b>twlgt</b> rA,SIMM	<b>tw 1</b> ,rA,SIMM	Trap if logically greater than
<b>twlgti</b> rA,SIMM	<b>twi 1</b> ,rA,SIMM	Trap immediate if logically greater than
<b>twlle</b> rA,SIMM	<b>tw 6</b> ,rA,SIMM	Trap if logically less than or equal
<b>twllei</b> rA,SIMM	<b>twi 6</b> ,rA,SIMM	Trap immediate if logically less than or equal
<b>twllt</b> rA,SIMM	<b>tw 2</b> ,rA,SIMM	Trap if logically less than
<b>twllti</b> rA,SIMM	<b>twi 2</b> ,rA,SIMM	Trap immediate if logically less than
<b>twlng</b> rA,SIMM	<b>tw 6</b> ,rA,SIMM	Trap if logically not greater than
<b>twlngi</b> rA,SIMM	<b>twi 6</b> ,rA,SIMM	Trap immediate if logically not greater than
<b>twlnl</b> rA,SIMM	<b>tw 5</b> ,rA,SIMM	Trap if logically not less than
<b>twlnli</b> rA,SIMM	<b>twi 5</b> ,rA,SIMM	Trap immediate if logically not less than
<b>twlt</b> rA,SIMM	<b>tw 16</b> ,rA,SIMM	Trap if less than
<b>twlti</b> rA,SIMM	<b>twi 16</b> ,rA,SIMM	Trap immediate if less than
<b>twne</b> rA,SIMM	<b>tw 24</b> ,rA,SIMM	Trap if not equal
<b>twnei</b> rA,SIMM	<b>twi 24</b> ,rA,SIMM	Trap immediate if not equal
<b>twng</b> rA,SIMM	<b>tw 20</b> ,rA,SIMM	Trap if not greater than
<b>twngi</b> rA,SIMM	<b>twi 20</b> ,rA,SIMM	Trap immediate if not greater than
<b>twnl</b> rA,SIMM	<b>tw 12</b> ,rA,SIMM	Trap if not less than
<b>twnli</b> rA,SIMM	<b>twi 12</b> ,rA,SIMM	Trap immediate if not less than

1. Simplified mnemonics for branch instructions that do not test a CR bit should not specify one; a programming error may occur.
2. The value in the BI operand selects CR $n$ [2], the EQ bit.
3. Instructions for which B0 is either 12 (branch if condition true) or 4 (branch if condition false) do not depend on the CTR value and can be alternately coded by incorporating the condition specified by the BI field, as described in [B.4.6: Simplified mnemonics that incorporate CR conditions \(eliminates BO and replaces BI with crS\)](#).
4. The value in the BI operand selects CR $n$ [0], the LT bit.
5. The value in the BI operand selects CR $n$ [1], the GT bit.
6. The value in the BI operand selects CR $n$ [3], the SO bit.

## Appendix C Programming examples

This appendix gives examples of how memory synchronization instructions can be used to emulate various synchronization primitives and to provide more complex forms of synchronization. It also describes multiple precision shifts.

### C.1 Synchronization

Examples in this appendix have a common form. After possible initialization, a conditional sequence begins with a load and reserve instruction that may be followed by memory accesses and computations that include neither a load and reserve nor a store conditional. The sequence ends with a store conditional with the same target address as the initial load and reserve. In most of the examples, failure of the store conditional causes a branch back to the load and reserve for a repeated attempt. On the assumption that contention is low, the conditional branch in the examples is optimized for the case in which the store conditional succeeds, by setting the branch-prediction bit appropriately. These examples focus on techniques for the correct modification of shared memory locations: see note <sup>(4)</sup> in [C.1.4: Synchronization notes](#), for a discussion of how the retry strategy can affect performance.

Load and reserve and store conditional instructions depend on the coherence mechanism of the system. Stores to a given location are coherent if they are serialized in some order, and no processor is able to observe a subset of those stores as occurring in a conflicting order.

Each load operation, whether ordinary or load and reserve, returns a value that has a well-defined source. The source can be the store or store conditional instruction that wrote the value, an operation by some other mechanism that accesses memory (for example, an I/O device), or the initial state of memory.

The function of an atomic read/modify/write operation is to read a location and write its next value, possibly as a function of its current value, all as a single atomic operation. We assume that locations accessed by read/modify/write operations are accessed coherently, so the concept of a value being the next in the sequence of values for a location is well defined. The conditional sequence, as defined above, provides the effect of an atomic read/modify/write operation, but not with a single atomic instruction. Let **addr** be the location that is the common target of the load and reserve and store conditional instructions. Then the guarantee the architecture makes for the successful execution of the conditional sequence is that no store into **addr** by another processor or mechanism has intervened between the source of the load and reserve and the store conditional.

For each of these examples, it is assumed that a similar sequence of instructions is used by all processes requiring synchronization on the accessed data.

*Note:* Because memory synchronization instructions have implementation dependencies (for example, the granularity at which reservations are managed), they must be used with care. The operating system should provide system library programs that use these instructions to implement the high-level synchronization functions (such as, test and set or compare and swap) needed by application programs. Application programs should use these library programs, rather than use memory synchronization instructions directly.

#### C.1.1 Synchronization primitives

The following examples show how the **lwarx** and **stwcx** instructions can be used to implement various synchronization primitives.

The sequences used to emulate the various primitives consist primarily of a loop using **lwarx** and **stwcx.** No additional synchronization is necessary, because the **stwcx.** will fail, clearing EQ, if the word loaded by **lwarx** has changed before the **stwcx.** is executed: see [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#) for details.

### Fetch and No-op

The fetch and no-op primitive atomically loads the current value in a word in memory.

In this example it is assumed that the address of the word to be loaded is in GPR3 and the data loaded are returned in GPR4.

```

loop:      lwarx      r4,0,r3      #load and reserve
           stwcx.    r4,0,r3      #store old value if still reserved
           bc        4,2,loop     #loop if lost reservation

```

If the **stwcx.** succeeds, it stores to the target location the same value that was loaded by the preceding **lwarx.** While the store is redundant with respect to the value in the location, its success ensures that the value loaded by the **lwarx** was the current value, that is, that the source of the value loaded by the **lwarx** was the last store to the location that preceded the **stwcx.** in the coherence order for the location.

### Fetch and store

The fetch and store primitive atomically loads and replaces a word in memory. In this example it is assumed that the address of the word to be loaded and replaced is in GPR3, the new value is in GPR4, and the old value is returned in GPR5.

```

loop:      lwarx      r5,0,r3      #load and reserve
           stwcx.    r4,0,r3      #store new value if still reserved
           bc        4,2,loop     #loop if lost reservation

```

### Fetch and add

The fetch and add primitive atomically increments a word in memory. In this example it is assumed that the address of the word to be incremented is in GPR3, the increment is in GPR4, and the old value is returned in GPR5.

```

loop:      lwarx      r5,0,r3      #load and reserve
           add       r0,r4,r5      #increment word
           stwcx.    r0,0,r3      #store new value if still reserved
           bc        4,2,loop     #loop if lost reservation

```

### Fetch and AND

The Fetch and AND primitive atomically ANDs a value into a word in memory.

In this example it is assumed that the address of the word to be ANDed is in GPR3, the value to AND into it is in GPR4, and the old value is returned in GPR5.

```

loop:      lwarx      r5,0,r3      #load and reserve
           and       r0,r4,r5      #AND word
           stwcx.    r0,0,r3      #store new value if still reserved
           bc        4,2,loop     #loop if lost reservation

```

This sequence can be changed to perform another Boolean operation atomically on a word in memory by changing the **and** to the desired Boolean instruction (**or**, **xor**, etc.).

### Test and set

This version of the test and set primitive atomically loads a word from memory, sets the word in memory to a nonzero value if the value loaded is zero, and sets the EQ bit of CR Field 0 to indicate whether the value loaded is zero.

In this example it is assumed that the address of the word to be tested is in GPR3, the new value (nonzero) is in GPR4, and the old value is returned in GPR5.

```

loop:      lwarx      r5,0,r3      #load and reserve
           cmpwi     r5,0         #done if word
           bc        4,2,done     #not equal to 0
           stwcx.   r4,0,r3      #try to store non-0
           bc        4,2,loop     #loop if lost reservation

done:

```

### Compare and swap

The compare and swap primitive atomically compares a value in a register with a word in memory, if they are equal stores the value from a second register into the word in memory, if they are unequal loads the word from memory into the first register, and sets CR0[EQ] to indicate the result of the comparison.

In this example it is assumed that the address of the word to be tested is in GPR3, the comparand is in GPR4 and the old value is returned there, and the new value is in GPR5.

```

loop:      lwarx      r6,0,r3      #load and reserve
           cmpw     r4,r6         #1st 2 operands equal?
           bc        4,2,exit     #skip if not
           stwcx.   r5,0,r3      #store new value if still reserved
           bc        4,2,loop     #loop if lost reservation

exit:     or         r4,r6,r6     #return value from memory

```

*Note: The semantics given for compare and swap above are based on those of the IBM System/370 compare and swap instruction. Other architectures may define a compare and swap instruction differently.*

*Compare and swap is shown primarily for pedagogical reasons. It is useful on machines that lack the better synchronization facilities provided by **lwarx** and **stwcx**. A major weakness of a System/370-style compare and swap instruction is that, although the instruction itself is atomic, it checks only that the old and current values of the word being tested are equal, with the result that programs that use such a compare and swap to control a shared resource can err if the word has been modified and the old value subsequently restored. The sequence shown above has the same weakness.*

*In some applications the second **bc** and/or the **or** can be omitted. The **bc** is needed only if the application requires that if CR0[EQ] on exit indicates not equal then GPR4 and GPR6 are not equal. The **or** is needed only if the application requires that if the comparands are not equal then the word from memory is loaded into the register with which it was compared (rather than into a third register). If any of these instructions is omitted, the resulting compare and swap does not obey System/370 semantics.*

## C.1.2 Lock acquisition and release

This example gives an algorithm for locking that demonstrates the use of synchronization with an atomic read/modify/write operation. A shared memory location, the address of which is an argument of the lock and unlock procedures, given by GPR3, is used as a lock, to control access to some shared resource such as a shared data structure. The lock is open

when its value is 0 and closed (locked) when its value is 1. Before accessing the shared resource the program executes the lock procedure, which sets the lock by changing its value from 0 to 1. To do this, the lock procedure calls `test_and_set`, which executes the code sequence shown in the test and set example of [C.1.1: Synchronization primitives](#), thereby atomically loading the old value of the lock, writing to the lock the new value (1) given in GPR4, returning the old value in GPR5 (not used below), and setting the EQ bit of CR Field 0 according to whether the value loaded is 0. The lock procedure repeats the `test_and_set` until it succeeds in changing the value of the lock from 0 to 1.

Because the shared resource must not be accessed until the lock has been set, the lock procedure contains an **isync** after the **bc** that checks for the success of `test_and_set`. The **isync** delays all subsequent instructions until all preceding instructions have completed.

```
lock:      mfspr      r6,LR      #save Link Register
          addi      r4,r0,1    #obtain lock:
loop:     bl       test_and_set # test-and-set
          bc       4,2,loop    # retry til old = 0
# Delay subsequent instructions til prior instructions finish
          isync
          mtspr    LR,r6      #restore Link Register
          blr
          #return
```

The unlock procedure stores a 0 to the lock location. Most applications that use locking require, for correctness, that if the access to the shared resource includes stores, the program must execute an **msync** before releasing the lock. The **msync** ensures that the program's modifications are performed with respect to other processors before the store that releases the lock is performed with respect to those processors. In this example, the unlock procedure begins with an **msync** for this purpose.

```
unlock:   msync                                #order prior stores
          addi      r1,r0,0    #before lock release
          stw      r1,0(r3)    #store 0 to lock location
          blr
          #return
```

### C.1.3 List insertion

This example shows how **lwarx** and **stwcx** can be used to implement simple insertion into a singly linked list. (Complicated list insertion, in which multiple values must be changed atomically, or in which the correct order of insertion depends on the contents of the elements, cannot be implemented in the manner shown below and requires a more complicated strategy such as using locks.)

The next element pointer from the list element after which the new element is to be inserted, here called the parent element, is stored into the new element, so that the new element points to the next element in the list: this store is performed unconditionally. Then the address of the new element is conditionally stored into the parent element, thereby adding the new element to the list.

In this example it is assumed that the address of the parent element is in GPR3, the address of the new element is in GPR4, and the next element pointer is at offset 0 from the start of the element. It is also assumed that the next element pointer of each list element is in a reservation granule separate from that of the next element pointer of all other list elements: see [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx](#).

```
loop:     lwarx    r2,0,r3    #get next pointer
          stw     r2,0(r4)    #store in new element
          msync   #order stw before stwcx.(can omit if not MP)
```

```

        stwcx.    r4,0,r3    #add new element to list
        bc       4,2,loop   #loop if stwcx. failed

```

In the preceding example, if two list elements have next element pointers in the same reservation granule then, in a multiprocessor, livelock can occur. (Livelock is a state in which processors interact in a way such that no processor makes progress.)

If list elements cannot be allocated such that each element's next element pointer is in a different reservation granule, livelock can be avoided with this more complicated sequence:

```

loop1:   lwz      r2,0(r3)    #get next pointer
         or       r5,r2,r2    #keep a copy
         stw     r2,0(r4)    #store in new element
         msync                    #order stw before stwcx.
loop2:   lwarx   r2,0,r3     #get it again
         cmpw   r2,r5       #loop if changed (someone
         bc    4,2,loop1    # else progressed)
         stwcx. r4,0,r3     #add new element to list
         bc    4,2,loop     #loop if failed

```

### C.1.4 Synchronization notes

1. In general, **lwarx** and **stwcx.** should be paired, with the same effective address used for both. The only exception is that an unpaired **stwcx.** to any (scratch) effective address can be used to clear any reservation held by the processor.
2. It is acceptable to execute a **lwarx** for which no **stwcx.** is executed. For example, this occurs in the test and set sequence shown above if the value loaded is not zero.
3. To increase the likelihood that forward progress is made, it is important that looping on **lwarx/stwcx.** pairs be minimized. For example, in the sequence shown above for test and set, this is achieved by testing the old value before attempting the store: were the order reversed, more **stwcx.** instructions might be executed, and reservations might more often be lost between the **lwarx** and the **stwcx.**
4. The manner in which **lwarx** and **stwcx.** are communicated to other processors and mechanisms, and between levels of the memory subsystem within a given processor is implementation-dependent (see [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#)). In some implementations performance may be improved by minimizing looping on a **lwarx** instruction that fails to return a desired value. For example, in the test and set example shown above, to stay in the loop until the word loaded is zero, `bne- $+12` can be changed to `bne- loop`. However, in some implementations better performance may be obtained by using an ordinary load instruction to do the initial checking of the value, as follows.

```

loop:    lwz      r5,0(r3)    #load the word
         cmpi   cr0,0,r5,0    #loop back if word
         bc    4,2,loop      # not equal to 0
         lwarx  r5,0,r3     #try again, reserving
         cmpi   cr0,0,r5,0    # (likely to succeed)
         bc    4,2,loop      #
         stwcx. r4,0,r3     #try to store non-0
         bc    4,2,loop      #loop if lost reservation

```

5. In a multiprocessor, livelock is possible if a loop containing a **lwarx/stwcx.** pair also contains an ordinary store instruction for which any byte of the affected memory area is in the reservation granule: see [Section 4.3.1.16: Atomic update primitives using lwarx and stwcx.](#) For example, the first code sequence shown in [C.1.3: List insertion](#), can

cause livelock if two list elements have next element pointers in the same reservation granule.

## C.2 Multiple-precision shifts

This section gives examples of how multiple-precision shifts can be programmed.

A multiple-precision shift is defined to be a shift of an N-word quantity (32-bit implementations), where  $N > 1$ . The quantity to be shifted is contained in N registers. The shift amount is specified either by an immediate value in the instruction or by a value in a register.

The examples shown below distinguish between the cases  $N=2$  and  $N > 2$ . If  $N=2$ , the shift amount may be in the range 0–63, which are the maximum ranges supported by the *Shift* instructions used. However if  $N > 2$ , the shift amount must be in the range 0–31 for the examples to yield the desired result. The specific instance shown for  $N > 2$  is  $N=3$ : extending those code sequences to larger N is straightforward, as is reducing them to the case  $N=2$  when the more stringent restriction on shift amount is met. For shifts with immediate shift amounts only the case  $N=3$  is shown, because the more stringent restriction on shift amount is always met.

In the examples it is assumed that GPRs 2 and 3 (and 4) contain the quantity to be shifted, and that the result is to be placed into the same registers. In all cases, for both input and result, the lowest-numbered register contains the highest-order part of the data and highest-numbered register contains the lowest-order part. For non-immediate shifts, the shift amount is assumed to be in GPR6. For immediate shifts, the shift amount is assumed to be greater than 0. GPRs 0 and 31 are used as scratch registers.

For  $N > 2$ , the number of instructions required is  $2N-1$  (immediate shifts) or  $3N-1$  (non-immediate shifts).

**Table 304. Shifts**

Left shifts	Right shifts
Shift Left Immediate, $N=3$ (shift amount < 32) rlwinm r2,r2,sh,0,31-sh rlwimi r2,r3,sh,32-sh,31 rlwinm r3,r3,sh,0,31-sh rlwimi r3,r4,sh,32-sh,31 rlwinm r4,r4,sh,0,31-sh	Shift Right Immediate, $N=3$ (shift amount < 32) rlwinm r4,r4,32-sh,sh,31 rlwimi r4,r3,32-sh,0,sh-1 rlwinm r3,r3,32-sh,sh,31 rlwimi r3,r2,32-sh,0,sh-1 rlwinm r2,r2,32-sh,sh,31
Shift Left, $N=2$ (shift amount < 64) subfic r31,r6,32 slw r2,r2,r6 srw r0,r3,r31 or r2,r2,r0 addi r31,r6,-32 slw r0,r3,r31 or r2,r2,r0 slw r3,r3,r6	Shift Right, $N=2$ (shift amount < 64) subfic r31,r6,32 srw r3,r3,r6 slw r0,r2,r31 or r3,r3,r0 addi r31,r6,-32 srw r0,r2,r31 or r3,r3,r0 srw r2,r2,r6



**Table 304. Shifts (continued)**

Left shifts	Right shifts
Shift Left, N=3 (shift amount < 32) subfic r31,r6,32 slw r2,r2,r6 srw r0,r3,r31 or r2,r2,r0 slw r3,r3,r6 srw r0,r4,r31 or r3,r3,r0 slw r4,r4,r6	Shift Right, N=3 (shift amount < 32) subfic r31,r6,32 srw r4,r4,r6 slw r0,r3,r31 or r4,r4,r0 srw r3,r3,r6 slw r0,r2,r31 or r3,r3,r0 srw r2,r2,r6
	Shift Right Algebraic Immediate, N=3 (shift amt < 32) rlwinm r4,r4,32-sh,sh,31 rlwimi r4,r3,32-sh,0,sh-1 rlwinm r3,r3,32-sh,sh,31 rlwimi r3,r2,32-sh,0,sh-1 srawi r2,r2,sh
	Shift Right Algebraic, N=2 (shift amount < 64) subfic r31,r6,32 srw r3,r3,r6 slw r0,r2,r31 or r3,r3,r0 addic. r31,r6,-32 sraw r0,r2,r31 bc 4,1,\$+8 ori r3,r0,0 sraw r2,r2,r6
	Shift Right Algebraic, N=3 (shift amount < 32) subfic r31,r6,32 srw r4,r4,r6 slw r0,r3,r31 or r4,r4,r0 srw r3,r3,r6 slw r0,r2,r31 or r3,r3,r0 sraw r2,r2,r6

### C.3 Floating point conversions

This section gives examples of how floating-point conversion instructions can be used to perform various conversions.

*Note:* Some of the examples use the optional *fsel* instruction. Care must be taken in using *fsel* if IEEE compatibility is required, or if the values being tested can be NaNs or infinities.

### C.3.1 Conversion from floating-point number to signed integer word

The full convert to signed integer word function can be implemented with the sequence shown below, assuming the floating-point value to be converted is in FPR1, the result is returned in GPR3, and a double word at displacement 'disp' from the address in GPR1 can be used as scratch space.

fctiw[z]	f2,f1	#convert to integer
stfd	f2,disp(r1)	#store float
lwa	r3,disp+4(r1)	#load word algebraic
		#(use lwz on a 32-bit implementation)

### C.3.2 Conversion from floating-point number to unsigned integer word

In a 32-bit implementation

The full convert to unsigned integer word function can be implemented with the sequence shown below, assuming the floating-point value to be converted is in FPR1, the value 0 is in FPR0, the value  $2^{32}-1$  is in FPR3, the value  $2^{31}$  is in FPR4, the result is returned in GPR3, and a double word at displacement 'disp' from the address in GPR1 can be used as scratch space.

fsel	f2,f1,f1,f0	#use 0 if < 0
fsub	f5,f3,f1	#use max if > max
fsel	f2,f5,f2,f3	
fsub	f5,f2,f4	#subtract $2^{31}$
fcmpu	cr2,f2,f4	#use diff if $\geq 2^{31}$
fsel	f2,f5,f5,f2	
fctiw[z]	f2,f2	#convert to integer
stfd	f2,disp(r1)	#store float
lwz	r3,disp+4(r1)	#load word
bc	12,8,\$+8	#add $2^{31}$ if input
xoris	r3,r3,0x8000	# was $\geq 2^{31}$

## C.4 Floating point selection

This section gives examples of how the optional floating select instruction (fsel) can be used to implement floating-point minimum and maximum functions, and certain simple forms of if-then-else constructions, without branching.

The examples show program fragments in an imaginary, C-like, high-level programming language, and the corresponding program fragment using fsel and other Book E instructions. In the examples, a, b, x, y, and z are floating-point variables, which are assumed to be in FPRs fa, fb, fx, fy, and fz. FPR fs is assumed to be available for scratch space.

---

**Warning:** Care must be taken in using fsel if IEEE compatibility is required, or if the values being tested can be NaNs or infinities: see [C.4.1: Notes](#).

---

Table 305. Comparison to zero

High-Level Language:	Book E:	Notes
if a $\leq$ 0.0 then x = y		
else x = z	fsel fx,fa,fy,fz	(1)
if a > 0.0 then x = y		
else x = z	fneg fs,fa	
fsel fx,fs,fz,fy	(1,2)	
if a = 0.0 then x = y		
else x = z	fsel fx,fa,fy,fz	
fneg fs,fa		
fsel fx,fs,fx,fz	(1)	

Table 306. Minimum and maximum

High-Level Language:	Book E:	Notes
x = min(a,b)	fsub fs,fa,fb	
fsel fx,fs,fb,fa	(3,4,5)	
x = max(a,b)	fsub fs,fa,fb	
fsel fx,fs,fa,fb	(3,4,5)	

Table 307. Simple if-then-else constructions

High-Level Language:	Book E:	Notes
if a $\leq$ b then x = y		
else x = z	fsub fs,fa,fb	
fsel fx,fs,fy,fz	(4,5)	
if a > b then x = y		
else x = z	fsub fs,fb,fa	
fsel fx,fs,fz,fy	(3,4,5)	
if a = b then x = y		
else x = z	fsub fs,fa,fb	
fsel fx,fs,fy,fz		
fneg fs,fs		
fsel fx,fs,fx,fz	(4,5)	

**C.4.1 Notes**

The following notes apply to the preceding examples and to the corresponding cases using the other three arithmetic relations (<, ≤, and ≠). They should also be considered when any other use of **fsel** is contemplated.

In these notes, the optimized program is the Book E program shown, and the unoptimized program (not shown) is the corresponding Book E program that uses **fcmpu** and branch conditional instructions instead of **fsel**.

1. The unoptimized program affects FPSCR[VXSNAN] and therefore may cause the system error handler to be invoked if the corresponding exception is enabled; the optimized program does not affect this bit. This property of the optimized program is incompatible with the IEEE standard.
2. The optimized program gives the incorrect result if  $a$  is a NaN.
3. The optimized program gives the incorrect result if  $a$  and/or  $b$  is a NaN (except that it may give the correct result in some cases for the minimum and maximum functions, depending on how those functions are defined to operate on NaNs).
4. The optimized program gives the incorrect result if  $a$  and  $b$  are infinities of the same sign. (Here it is assumed that invalid operation exceptions are disabled, in which case the result of the subtraction is a NaN. The analysis is more complicated if invalid operation exceptions are enabled, because in that case the target register of the subtraction is unchanged.)
5. The optimized program affects FPSCR[OX, UX, XX, VXISI], and therefore may cause the system error handler to be invoked if the corresponding exceptions are enabled; the unoptimized program does not affect these bits. This property of the optimized program is incompatible with the IEEE standard.

## Appendix D Guidelines for 32-bit book E

This appendix provides guidelines used by 32-bit Book E implementations; a set of guidelines is also outlined for software developers. Application software written to these guidelines can be labeled 32-bit Book E applications and can be expected to execute properly on all implementations of Book E, both 32-bit and 64-bit implementations.

32-bit Book E implementations execute applications that adhere to the software guidelines for 32-bit Book E software outlined in this appendix and are not expected to properly execute 64-bit Book E applications or any applications not adhering to these guidelines (that is, 64-bit Book E applications).

### D.1 Registers on 32-bit book E implementations

Book E defines 32- and 64-bit registers. All 32-bit registers are supported as defined in Book E. However, except for the 64-bit FPRs, only bits 32–63 of Book E's 64-bit registers are required to be implemented in hardware in 32-bit Book E implementation. Such 64-bit registers include LR, CTR, 32 GPRs, SRR0, and CSRR0. Book E makes no restrictions regarding implementing a subset of the 64-bit floating-point architecture.

Likewise, other than floating-point instructions, all instructions defined to return a 64-bit result return only bits 32–63 of the result on a 32-bit Book E implementation.

### D.2 Addressing on 32-bit book E implementations

Only bits 32–63 of the 64-bit Book E instruction and data memory effective addresses need to be calculated and presented to main memory, so a 32-bit implementation can bypass prepending the 32 zeros when implementing these instructions. For branch to LR and branch to CR instructions, given that LR and CTR are implemented as 32-bit registers, only 2 zeros need to be concatenated to the right of bits 32–61 of these registers to form the 32-bit branch target address.

The simplest implementation of next sequential instruction address computation suggests allowing effective address computations to wrap from 0xFFFF\_FFFC to 0x0000\_0000. This wrapping is required of PowerPC implementations. For 32-bit Book E applications, there appears little if any benefit to allowing this wrapping behavior. Book E specifies that the situation where the computation of the next sequential instruction address after address 0xFFFF\_FFFC is undefined. (Note that the next sequential instruction address after address 0xFFFF\_FFFC on a 64-bit Book E implementation is 0x0000\_0001\_0000\_0000.)

### D.3 TLB fields on 32-bit book E implementations

32-bit Book E implementations should support bits 32–53 of the effective page number (EPN) field in the TLB. This size provides support for a 32-bit effective address, which PowerPC ABIs may have come to expect to be available. 32-bit Book E implementations may support greater than 32-bit real addresses by supporting more than bits 32–53 of the real page number (RPN) field in the TLB.

## D.4 32-bit book E software guidelines

### D.4.1 32-bit instruction selection

Generally speaking, 32-bit software should avoid instructions that depend on any particular setting of bits 0–31 of any 64-bit application-accessible system register, including GPRs, for producing the correct 32-bit results. Context switching is not required to preserve the upper 32 bits of application-accessible 64-bit system registers and insertion of arbitrary settings of those upper 32 bits at arbitrary times during the execution of the 32-bit application must not affect the final result.

### D.4.2 32-bit addressing

Book E provides a complete set of data memory access instructions that perform a modulo  $2^{32}$  on the computed effective address and then prepend 32 zeros to produce the full 64-bit address. Book E also provides a complete set of branch instructions that perform a modulo  $2^{32}$  on the computed branch target effective address and then prepend 32 zeros to produce the full 64-bit branch target address. On a 32-bit Book E implementation, these instructions are executed as defined, but without prepending the 32 zeros (only the low-order 32 bits of the address are calculated). On a 64-bit implementation, executing these instructions as defined provides the effect of restricting the application to the lowest 32-bit address space.

However, there is one exception. Next sequential instruction address computations (not a taken branch) are not defined for 32-bit Book E applications when the current instruction address is `0xFFFF_FFFC`. On a 32-bit Book E implementation, the instruction address could simply wrap to `0x0000_0000`, providing the same effect that is required in the PowerPC Architecture. However, when the 32-bit Book E application is executed on a 64-bit Book E implementation, the next sequential instruction address calculated will be `0x0000_0001_0000_0000` and not `0x0000_0000_0000_0000`. To avoid this problem the 32-bit Book E application must either avoid this situation by not allowing code to span this address boundary, or requiring a branch absolute to address 0 be placed at address `0xFFFF_FFFC` to emulate the wrap. Either of these approaches allows the application to execute on 32-bit and 64-bit Book E implementations.

## Appendix E Embedded floating-point results

This appendix summarizes results of various types of floating-point operations on various combinations of input operands. Flag settings are performed on appropriate element flags.

### E.1 Notation conventions and general rules

For all tables in this appendix, the annotation and general rules in [Table 308](#) apply.

**Table 308. Notation conventions and general rules**

Notation	Description
*	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 0x7F7_FFFFF. The encoding for double-precision is 0x7FEF_FFFF_FFFF_FFFF.
nmax	Denotes the maximum normalized negative number. The encoding for single-precision is 0xFF7F_FFFF. The encoding for double-precision is 0xFFEF_FFFF_FFFF_FFFF.
pmin	Denotes the minimum normalized positive number. The encoding for single-precision is 0x00800000. The encoding for double-precision is 0x0010_0000_0000_0000.
nmin	Denotes the minimum normalized negative number. The encoding for single-precision is 0x8080_0000. The encoding for double-precision is 0x8010_0000_0000_0000.
Calculations that overflow or underflow saturate.	Overflow for operations that have a floating-point result force the result to <i>max</i> . 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 0x7FFF_FFFF (positive) or 0x8000_0000 (negative). Overflow for operations that have an unsigned integer result force the result to 0xFFFF_FFFF (positive) or 0x0000_0000 (negative).
1 (superscript)	Denotes that the sign of the result is positive when the signs of Operand A and 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 signs of Operand A and 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



## E.2 Add, subtract, multiply, and divide results

Table 309 lists results for add, subtract, multiply, and divide operations.

**Table 309. Floating-point results summary—add, sub, mul, div**

Operation	Operand A	Operand B	Result	FINV	FOVF	FUNF	FDBZ	FINX
<b>Add</b>								
Add	$\infty$	$\infty$	amax	1	0	0	0	0
Add	$\infty$	NaN	amax	1	0	0	0	0
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 <sup>1</sup>	1	0	0	0	0
Add	denorm	zero	zero <sup>1</sup>	1	0	0	0	0
Add	denorm	norm	operand_b <sup>4</sup>	1	0	0	0	0
Add	zero	$\infty$	bmax	1	0	0	0	0
Add	zero	NaN	bmax	1	0	0	0	0
Add	zero	denorm	zero <sup>1</sup>	1	0	0	0	0
Add	zero	zero	zero <sup>1</sup>	0	0	0	0	0
Add	zero	norm	operand_b <sup>4</sup>	0	0	0	0	0
Add	norm	$\infty$	bmax	1	0	0	0	0
Add	norm	NaN	bmax	1	0	0	0	0
Add	norm	denorm	operand_a <sup>4</sup>	1	0	0	0	0
Add	norm	zero	operand_a <sup>4</sup>	0	0	0	0	0
Add	norm	norm	_Calc_	0	*	*	0	*
<b>Subtract</b>								
Sub	$\infty$	$\infty$	amax	1	0	0	0	0
Sub	$\infty$	NaN	amax	1	0	0	0	0
Sub	$\infty$	denorm	amax	1	0	0	0	0
Sub	$\infty$	zero	amax	1	0	0	0	0

Table 309. Floating-point results summary—add, sub, mul, div (continued)

Operation	Operand A	Operand B	Result	FINV	FOVF	FUNF	FDBZ	FINX
Sub	$\infty$	Norm	amax	1	0	0	0	0
Sub	NaN	$\infty$	amax	1	0	0	0	0
Sub	NaN	NaN	amax	1	0	0	0	0
Sub	NaN	denorm	amax	1	0	0	0	0
Sub	NaN	zero	amax	1	0	0	0	0
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 <sup>2</sup>	1	0	0	0	0
Sub	denorm	zero	zero <sup>2</sup>	1	0	0	0	0
Sub	denorm	norm	-operand_b <sup>4</sup>	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 <sup>2</sup>	1	0	0	0	0
Sub	zero	zero	zero <sup>2</sup>	0	0	0	0	0
Sub	zero	norm	-operand_b <sup>4</sup>	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 <sup>4</sup>	1	0	0	0	0
Sub	norm	zero	operand_a <sup>4</sup>	0	0	0	0	0
Sub	norm	norm	_Calc_	0	*	*	0	*
Multiply <sup>3</sup>								
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

Table 309. Floating-point results summary—add, sub, mul, div (continued)

Operation	Operand A	Operand B	Result	FINV	FOVF	FUNF	FDBZ	FINX
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
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 <sup>3</sup>								
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

**Table 309. Floating-point results summary—add, sub, mul, div (continued)**

Operation	Operand A	Operand B	Result	FINV	FOVF	FUNF	FDBZ	FINX
Div	norm	denorm	max	1	0	0	0	0
Div	norm	zero	max	0	0	0	1	0
Div	norm	norm	_Calc_	0	*	*	0	*

### E.3 Double- to single-precision conversion

*Table 310* lists results for double- to single-precision conversion.

**Table 310. Floating-point results summary—single convert from double**

Operand B	efscfd result	FINV	FOVF	FUNF	FDBZ	FINX
+∞	pmax	1	0	0	0	0
-∞	nmax	1	0	0	0	0
+NaN	pmax	1	0	0	0	0
-NaN	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	*

## E.4 Single- to double-precision conversion

[Table 311](#) lists results for single- to double-precision conversion.

**Table 311. 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
+NaN	pmax	1	0	0	0	0
-NaN	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

## E.5 Conversion to unsigned

[Table 312](#) lists results for conversion to unsigned operations.

**Table 312. 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	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	*

## E.6 Conversion to signed

Table 313 lists results for conversion to signed operations.

**Table 313. Floating-point results summary—convert to signed**

Operand B	Integer result ctsi[d][z]	Fractional result ctsf	FINV	FOVF	FUNF	FDBZ	FINX
+∞	0x7FFF_FFFF	0x7FFF_FFFF	1	0	0	0	0
-∞	0x8000_0000	0x8000_0000	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	*

## E.7 Conversion from unsigned

Table 314 lists results for conversion from unsigned operations.

**Table 314. Floating-point 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	*

## E.8 Conversion from signed

Table 315 lists results for conversion from signed operations.

**Table 315. 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	*

## E.9 \*abs, \*nabs, and \*neg operations

Table 316 lists results for \*abs, \*nabs, and \*neg operations.

**Table 316. Floating-point results summary—\*abs, \*nabs, \*neg**

Operand A	*abs	*nabs	*neg	FINV	FOVF	FUNF	FDBZ	FINX
+∞	pmax   +∞	nmax   -∞	-amax   -∞	1	0	0	0	0
-∞	pmax   +∞	nmax   -∞	-amax   +∞	1	0	0	0	0
+NaN	pmax   NaN	nmax   -NaN	-amax   -NaN	1	0	0	0	0
-NaN	pmax   NaN	nmax   -NaN	-amax   +NaN	1	0	0	0	0
+denorm	+zero   +denorm	-zero   -denorm	-zero   -denorm	1	0	0	0	0
-denorm	+zero   +denorm	-zero   -denorm	+zero   +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

## Glossary

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this book. Some of the terms and definitions included in the glossary are reprinted from IEEE Standard 754-1985, *IEEE Standard for Binary Floating-Point Arithmetic*, copyright ©1985 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE.

**Table 317. Glossary**

Name	Description
<b>A</b>	
<b>Architecture</b>	A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible <i>implementations</i> .
<b>Asynchronous interrupt</b>	<i>interrupts</i> that are caused by events external to the processor's execution. In this document, the term <i>asynchronous interrupt</i> is used interchangeably with the word <i>interrupt</i> .
<b>Atomic access</b>	A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The PowerPC architecture implements <i>atomic accesses</i> through the <b>lwarx/stwcx</b> . instruction pair.
<b>B</b>	
<b>Biased exponent</b>	An <i>exponent</i> whose range of values is shifted by a constant (bias). Typically a bias is provided to allow a range of positive values to express a range that includes both positive and negative values.
<b>Big-endian</b>	A byte-ordering method in memory where the address <i>n</i> of a word corresponds to the <i>most-significant byte</i> . In an addressed memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the <i>most-significant byte</i> . See <i>Little-endian</i> .
<b>Boundedly undefined</b>	<p>A characteristic of certain operation results that are not rigidly prescribed by the PowerPC architecture. Boundedly-undefined results for a given operation may vary among implementations and between execution attempts in the same implementation.</p> <p>Although the architecture does not prescribe the exact behavior for when results are allowed to be <i>boundedly undefined</i>, the results of executing instructions in contexts where results are allowed to be <i>boundedly undefined</i> are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.</p>
<b>Branch prediction</b>	The process of guessing whether a branch will be taken. Such predictions can be correct or incorrect; the term 'predicted' as it is used here does not imply that the prediction is correct (successful). The PowerPC architecture defines a means for <i>static branch</i> prediction as part of the instruction encoding.



**Table 317. Glossary (continued)**

Name	Description
<b>Branch resolution</b>	The determination of whether a branch is taken or not taken. A branch is said to be resolved when the processor can determine which instruction path to take. If the branch is resolved as predicted, the instructions following the predicted branch that may have been speculatively executed can complete (see <i>Completion</i> ). If the branch is not resolved as predicted, instructions on the mispredicted path, and any results of speculative execution, are purged from the pipeline and fetching continues from the nonpredicted path.
<b>C</b>	
<b>Cache</b>	High-speed memory containing recently accessed data or instructions (subset of main memory).
<b>Cache block</b>	A small region of contiguous memory that is copied from memory into a <i>cache</i> . The size of a <i>cache block</i> may vary among processors; the maximum block size is one <i>page</i> . In PowerPC processors, <i>cache coherency</i> is maintained on a cache-block basis. Note that the term <i>cache block</i> is often used interchangeably with 'cache line.'
<b>Cache coherency</b>	An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor's cache.
<b>Cache flush</b>	An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush ( <b>dcbf</b> ) instruction.
<b>Caching-inhibited</b>	A memory update policy in which the cache is bypassed and the load or store is performed to or from main memory.
<b>Cast out</b>	A <i>cache block</i> that must be written to memory when a cache miss causes a <i>cache block</i> to be replaced.
<b>Changed bit</b>	One of two <i>page history bits</i> found in each <i>page table entry</i> (PTE). The processor sets the changed bit if any store is performed into the <i>page</i> . See also <i>Page access history bits</i> and <i>Referenced bit</i> .
<b>Clean</b>	An operation that causes a <i>cache block</i> to be written to memory, if modified, and then left in a valid, unmodified state in the cache.
<b>Clear</b>	To cause a bit or bit field to register a value of zero. See also <i>Set</i> .
<b>Completion</b>	Completion occurs when an instruction has finished executing, written back any results, and is removed from the completion queue (CQ). When an instruction completes, it is guaranteed that this instruction and all previous instructions can cause no interrupts.
<b>Context synchronization.</b>	An operation that ensures that all instructions in execution complete past the point where they can produce an <i>interrupt</i> , that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are <i>fetched</i> and executed in the new context. Context synchronization may result from executing specific instructions (such as <b>isync</b> or <b>rfi</b> ) or when certain events occur (such as an <i>interrupt</i> ).
<b>D</b>	
<b>Denormalized number</b>	A nonzero floating-point number whose exponent has a reserved value, usually the format's minimum, and whose explicit or implicit leading significand bit is zero.
<b>E</b>	
<b>Effective address (EA)</b>	The 32-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a <i>physical memory</i> address or an I/O address.

Table 317. Glossary (continued)

Name	Description
<b>Exception</b>	A condition that, if enabled, generates an interrupt.
<b>Execution synchronization</b>	A mechanism by which all instructions in execution are architecturally complete before beginning execution (appearing to begin execution) of the next instruction. Similar to context synchronization but doesn't force the contents of the instruction buffers to be deleted and refetched.
<b>Exponent</b>	In the binary representation of a floating-point number, the exponent is the component that normally signifies the integer power to which the value two is raised in determining the value of the represented number. See also <i>Biased exponent</i> .
<b>F</b>	
<b>Fetch</b>	Instruction retrieval from either the cache or main memory and placing them into the instruction queue.
<b>Finish</b>	Finishing occurs in the last cycle of execution. In this cycle, the CQ entry is updated to indicate that the instruction has finished executing.
<b>Floating-point register (FPR)</b>	Any of the 32 registers in the floating-point register file. These registers provide the source operands and destination results for floating-point instructions. Load instructions move data from memory to FPRs and store instructions move data from FPRs to memory. The FPRs are 64 bits wide and store floating-point values in double-precision format.
<b>Floating-point unit</b>	The functional unit in a processor responsible for executing all floating-point instructions.
<b>Flush</b>	An operation that causes a cache block to be invalidated and the data, if modified, to be written to memory.
<b>Fraction</b>	In the binary representation of a floating-point number, the field of the <i>significand</i> that lies to the right of its implied binary point.
<b>G</b>	
<b>General-purpose register (GPR)</b>	Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.
<b>Guarded</b>	The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.
<b>H</b>	
<b>Harvard architecture</b>	An architectural model featuring separate caches and other memory management resources for instructions and data.
<b>I</b>	
<b>IEEE 754</b>	A standard written by the Institute of Electrical and Electronics Engineers that defines operations and representations of binary floating-point numbers.
<b>Illegal instructions</b>	A class of instructions that are not implemented for a particular PowerPC processor. These include instructions not defined by the PowerPC architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations instructions that are defined only for 32-bit implementations are considered to be illegal instructions.

Table 317. Glossary (continued)

Name	Description
<b>Implementation</b>	A particular processor that conforms to the PowerPC architecture, but may differ from other architecture-compliant implementations for example in design, feature set, and implementation of <i>optional</i> features. The PowerPC architecture has many different implementations.
<b>Imprecise interrupt</b>	A type of <i>synchronous interrupt</i> that is allowed not to adhere to the precise interrupt model (see <i>Precise interrupt</i> ). The PowerPC architecture allows only floating-point exceptions to be handled imprecisely.
<b>Integer unit.</b>	The functional unit responsible for executing all integer instructions.
<b>In order</b>	<b>An</b> aspect of an operation that adheres to a sequential model. An operation is said to be performed in-order if, at the time that it is performed, it is known to be required by the sequential execution model. See <i>Out-of-order</i> .
<b>Instruction latency</b>	The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.
<b>Interrupt</b>	A condition encountered by the processor that requires special, supervisor-level processing.
<b>Interrupt handler</b>	A software routine that executes when an interrupt is taken. Normally, the interrupt handler corrects the condition that caused the interrupt, or performs some other meaningful task (that may include aborting the program that caused the interrupt).
<b>K</b>	
<b>Kill</b>	An operation that causes a <i>cache block</i> to be invalidated without writing any modified data to memory.
<b>L</b>	
<b>Latency</b>	The number of clock cycles necessary to execute an instruction and make ready the results of that execution for a subsequent instruction.
<b>L2 cache</b>	See <i>Secondary cache</i> .
<b>Least-significant bit (lsb)</b>	The bit of least value in an address, register, field, data element, or instruction encoding.
<b>Least-significant byte (LSB)</b>	The byte of least value in an address, register, data element, or instruction encoding.
<b>Little-endian</b>	A byte-ordering method in memory where the address <i>n</i> of a word corresponds to the <i>least-significant byte</i> . In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the <i>most-significant byte</i> . See <i>Big-endian</i> .
<b>M</b>	
<b>Mantissa</b>	The decimal part of logarithm.
<b>Memory access ordering</b>	The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.
<b>Memory-mapped accesses</b>	Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.
<b>Memory coherency</b>	An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.
<b>Memory consistency</b>	Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).

Table 317. Glossary (continued)

Name	Description
<b>Memory management unit (MMU)</b>	The functional unit that is capable of translating an <i>effective</i> (logical) <i>address</i> to a physical address, providing protection mechanisms, and defining caching methods.
<b>Most-significant bit (msb)</b>	The highest-order bit in an address, registers, data element, or instruction encoding.
<b>Most-significant byte (MSB)</b>	The highest-order byte in an address, registers, data element, or instruction encoding.
<b>N</b>	
<b>NaN</b>	An abbreviation for not a number; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs and quiet NaNs.
<b>No-op</b>	No-operation. A single-cycle operation that does not affect registers or generate bus activity.
<b>Normalization</b>	A process by which a floating-point value is manipulated such that it can be represented in the format for the appropriate precision (single- or double-precision). For a floating-point value to be representable in the single- or double-precision format, the leading implied bit must be a 1.
<b>O</b>	
<b>OEA (operating environment architecture)</b>	The level of the architecture that describes PowerPC memory management model, supervisor-level registers, synchronization requirements, and the interrupt model. It also defines the time-base feature from a supervisor-level perspective. Implementations that conform to the PowerPC OEA also conform to the PowerPC UISA and VEA.
<b>Optional</b>	A feature, such as an instruction, a register, or an interrupt, that is defined by the PowerPC architecture but not required to be implemented.
<b>Out-of-order</b>	<b>An</b> aspect of an operation that allows it to be performed ahead of one that may have preceded it in the sequential model, for example, speculative operations. An operation is said to be performed out-of-order if, at the time that it is performed, it is not known to be required by the sequential execution model. See <i>In-order</i> .
<b>Out-of-order execution</b>	A technique that allows instructions to be issued and completed in an order that differs from their sequence in the instruction stream.
<b>Overflow</b>	An condition that occurs during arithmetic operations when the result cannot be stored accurately in the destination register(s). For example, if two 32-bit numbers are multiplied, the result may not be representable in 32 bits. Since 32-bit registers cannot represent this sum, an overflow condition occurs.
<b>P</b>	
<b>Page</b>	A region in memory. The OEA defines a page as a 4-Kbyte area of memory, aligned on a 4-Kbyte boundary.
<b>Page fault</b>	A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a <i>page</i> not currently resident in <i>physical memory</i> . On PowerPC processors, a page fault interrupt condition occurs when a matching, valid <i>page table entry</i> (PTE[V] = 1) cannot be located.
<b>Physical memory</b>	The actual memory that can be accessed through the system's memory bus.
<b>Pipelining</b>	A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

Table 317. Glossary (continued)

Name	Description
<b>Precise interrupts</b>	A category of interrupt for which the pipeline can be stopped so instructions that preceded the faulting instruction can complete and subsequent instructions can be flushed and redispached after interrupt handling has completed. See <i>Imprecise interrupts</i> .
<b>Primary opcode</b>	The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction.
<b>Program order</b>	The order of instructions in an executing program. More specifically, this term is used to refer to the original order in which program instructions are fetched into the instruction queue from the cache.
<b>Protection boundary</b>	A boundary between <i>protection domains</i> .
<b>Q</b>	
<b>Quiet NaN.</b>	A type of <i>NaN</i> that can propagate through most arithmetic operations without signaling interrupts. A quiet NaN is used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when invalid. See <i>Signaling NaN</i> .
<b>R</b>	
<b>Record bit</b>	Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.
<b>Referenced bit</b>	One of two <i>page history bits</i> found in each <i>page table entry</i> . The processor sets the <i>referenced bit</i> whenever the page is accessed for a read or write. See also <i>Page access history bits</i> .
<b>Register indirect addressing</b>	A form of addressing that specifies one GPR that contains the address for the load or store.
<b>Register indirect with immediate index addressing</b>	A form of addressing that specifies an immediate value to be added to the contents of a specified GPR to form the target address for the load or store.
<b>Register indirect with index addressing</b>	A form of addressing that specifies that the contents of two GPRs be added together to yield the target address for the load or store.
<b>Rename register</b>	Temporary buffers used by instructions that have finished execution but have not completed.
<b>Reservation</b>	The processor establishes a reservation on a <i>cache block</i> of memory space when it executes an <i>lwarx</i> instruction to read a memory semaphore into a GPR.
<b>Reservation station</b>	A buffer between the dispatch and execute stages that allows instructions to be dispatched even though the results of instructions on which the dispatched instruction may depend are not available.
<b>RISC (reduced instruction set computing)</b>	An <i>architecture</i> characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.
<b>S</b>	
<b>Secondary cache</b>	A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.
<b>Set (v)</b>	To write a nonzero value to a bit or bit field; the opposite of <i>clear</i> . The term 'set' may also be used to generally describe the updating of a bit or bit field.

Table 317. Glossary (continued)

Name	Description
<b>Set (n)</b>	A subdivision of a <i>cache</i> . Cacheable data can be stored in a given location in one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose <i>cache block</i> corresponding to that address was used least recently. See <i>Set-associative</i> .
<b>Set-associative</b>	Aspect of cache organization in which the cache space is divided into sections, called <i>sets</i> . The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.
<b>Signaling NaN</b>	A type of <i>NaN</i> that generates an invalid operation program interrupt when it is specified as arithmetic operands. See <i>Quiet NaN</i> .
<b>Significand</b>	The component of a binary floating-point number that consists of an explicit or implicit leading bit to the left of its implied binary point and a fraction field to the right.
<b>Simplified mnemonics</b>	Assembler mnemonics that represent a more complex form of a common operation.
<b>Snooping</b>	Monitoring addresses driven by a bus master to detect the need for coherency actions.
<b>Split-transaction</b>	A transaction with independent request and response tenures.
<b>Stall</b>	An occurrence when an instruction cannot proceed to the next stage.
<b>Static branch prediction</b>	Mechanism by which software (for example, compilers) can hint to the machine hardware about the direction a branch is likely to take.
<b>Superscalar</b>	A superscalar processor is one that can dispatch multiple instructions concurrently from a conventional linear instruction stream. In a superscalar implementation, multiple instructions can be in the same stage at the same time.
<b>Supervisor mode</b>	The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.
<b>Synchronization</b>	A process to ensure that operations occur strictly <i>in order</i> . See <i>Context synchronization</i> and <i>Execution synchronization</i> .
<b>Synchronous interrupt</b>	An <i>interrupt</i> that is generated by the execution of a particular instruction or instruction sequence. There are two types of synchronous interrupts, <i>precise</i> and <i>imprecise</i> .
<b>System memory</b>	The physical memory available to a processor.
<b>T</b>	
<b>TLB (translation lookaside buffer)</b>	A cache that holds recently-used <i>page table entries</i> .
<b>Throughput</b>	The measure of the number of instructions that are processed per clock cycle.
<b>U</b>	
<b>UISA (user instruction set architecture)</b>	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, floating-point memory conventions and interrupt model as seen by user programs, and the memory and programming models.
<b>Underflow</b>	A condition that occurs during arithmetic operations when the result cannot be represented accurately in the destination register. For example, underflow can happen if two floating-point fractions are multiplied and the result requires a smaller <i>exponent</i> and/or <i>mantissa</i> than the single-precision format can provide. In other words, the result is too small to be represented accurately.

Table 317. Glossary (continued)

Name	Description
<b>User mode</b>	The operating state of a processor used typically by application software. In user mode, software can access only certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.
<b>V</b>	
<b>VEA (virtual environment architecture)</b>	The level of the <i>architecture</i> that describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time-base facility from a user-level perspective. <i>Implementations</i> that conform to the PowerPC VEA also adhere to the UISA, but may not necessarily adhere to the OEA.
<b>Virtual address</b>	An intermediate address used in the translation of an <i>effective address</i> to a physical address.
<b>Virtual memory</b>	The address space created using the memory management facilities of the processor. Program access to <i>virtual memory</i> is possible only when it coinc
<b>W</b>	
<b>Way</b>	A location in the cache that holds a cache block, its tags and status bits.
<b>Word</b>	A 32-bit data element.
<b>Write-back</b>	A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is <i>cast out</i> to make room for newer data.
<b>Write-through</b>	A cache memory update policy in which all processor write cycles are written to both the cache and memory.

## Revision history

**Table 318. Document revision history**

Date	Revision	Changes
29-Nov-2007	1	Initial release.
25-May-2015	2	Editorial and formatting changes throughout reference manual. Updated Introduction. Updated Footnote in <a href="#">Table 269</a> , <a href="#">Table 270</a> , <a href="#">Table 271</a> , and <a href="#">Table 272</a> .



**IMPORTANT NOTICE – PLEASE READ CAREFULLY**

STMicroelectronics NV and its subsidiaries ("ST") reserve the right to make changes, corrections, enhancements, modifications, and improvements to ST products and/or to this document at any time without notice. Purchasers should obtain the latest relevant information on ST products before placing orders. ST products are sold pursuant to ST's terms and conditions of sale in place at the time of order acknowledgement.

Purchasers are solely responsible for the choice, selection, and use of ST products and ST assumes no liability for application assistance or the design of Purchasers' products.

No license, express or implied, to any intellectual property right is granted by ST herein.

Resale of ST products with provisions different from the information set forth herein shall void any warranty granted by ST for such product.

ST and the ST logo are trademarks of ST. All other product or service names are the property of their respective owners.

Information in this document supersedes and replaces information previously supplied in any prior versions of this document.

© 2015 STMicroelectronics – All rights reserved

