



---

PPC440x5 CPU Core

User's Manual

---

Preliminary

**SA14-2613-03**

July 15, 2003

***PowerPC***<sup>®</sup>



© Copyright International Business Machines Corporation 2003

All Rights Reserved  
Printed in the United States of America July 2003

The following are trademarks of International Business Machines Corporation in the United States, or other countries, or both.

IBM            IBM Logo  
CoreConnect  
PowerPC    PowerPC logo  
PowerPC Architecture  
RISCTrace   RISCWatch

Other company, product, and service names may be trademarks or service marks of others.

**The information contained in this document is subject to change or withdrawal at any time without notice and is being provided on an "AS IS" basis without warranty or indemnity of any kind, whether express or implied, including without limitation, the implied warranties of non-infringement, merchantability, or fitness for a particular purpose. Any products, services, or programs discussed in this document are sold or licensed under IBM's standard terms and conditions, copies of which may be obtained from your local IBM representative. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties.**

**Without limiting the generality of the foregoing, any performance data contained in this document was determined in a specific or controlled environment and not submitted to any formal IBM test. Therefore, the results obtained in other operating environments may vary significantly. Under no circumstances will IBM be liable for any damages whatsoever arising out of or resulting from any use of the document or the information contained herein.**

While the information contained herein is believed to be accurate, such information is preliminary, and should not be relied upon for accuracy or completeness, and no representations or warranties of accuracy or completeness are made.

**Note:** This document contains information on products in the sampling and/or initial production phases of development. This information is subject to change without notice. Verify with your IBM field applications engineer that you have the latest version of this document before finalizing a design.

IBM Microelectronics Division  
1580 Route 52, Bldg. 504  
Hopewell Junction, NY 12533-6351

The IBM home page can be found at <http://www.ibm.com>

The IBM Microelectronics Division home page can be found at <http://www.ibm.com/chips>

title.fm.  
July 15, 2003

# Contents

<b>Figures .....</b>	<b>15</b>
<b>Tables .....</b>	<b>19</b>
<b>About This Book .....</b>	<b>23</b>
<b>1. Overview .....</b>	<b>27</b>
PPC440x5 Features .....	27
The PPC440x5 as a PowerPC Implementation .....	29
PPC440x5 Organization .....	30
Superscalar Instruction Unit .....	30
Execution Pipelines .....	31
Instruction and Data Cache Controllers .....	31
Instruction Cache Controller (ICC) .....	31
Data Cache Controller (DCC) .....	32
Memory Management Unit (MMU) .....	32
Timers .....	34
Debug Facilities .....	34
Debug Modes .....	34
Development Tool Support .....	35
Core Interfaces .....	35
Processor Local Bus (PLB) .....	36
Device Control Register (DCR) Interface .....	36
Auxiliary Processor Unit (APU) Port .....	36
JTAG Port .....	37
<b>2. Programming Model .....</b>	<b>39</b>
Storage Addressing .....	39
Storage Operands .....	39
Effective Address Calculation .....	41
Data Storage Addressing Modes .....	41
Instruction Storage Addressing Modes .....	41
Byte Ordering .....	42
Structure Mapping Examples .....	43
Instruction Byte Ordering .....	44
Data Byte Ordering .....	45
Byte-Reverse Instructions .....	46
Registers .....	47
Register Types .....	51
General Purpose Registers .....	51
Special Purpose Registers .....	51
Condition Register .....	52
Machine State Register .....	52
Device Control Registers .....	52
Instruction Classes .....	52
Defined Instruction Class .....	53
Allocated Instruction Class .....	53
Preserved Instruction Class .....	54

Reserved Instruction Class .....	55
Implemented Instruction Set Summary .....	55
Integer Instructions .....	56
Integer Storage Access Instructions .....	56
Integer Arithmetic Instructions .....	57
Integer Logical Instructions .....	58
Integer Compare Instructions .....	58
Integer Trap Instructions .....	58
Integer Rotate Instructions .....	58
Integer Shift Instructions .....	59
Integer Select Instruction .....	59
Branch Instructions .....	59
Processor Control Instructions .....	59
Condition Register Logical Instructions .....	60
Register Management Instructions .....	60
System Linkage Instructions .....	60
Processor Synchronization Instruction .....	60
Storage Control Instructions .....	61
Cache Management Instructions .....	61
TLB Management Instructions .....	61
Storage Synchronization Instructions .....	61
Allocated Instructions .....	62
Branch Processing .....	62
Branch Addressing .....	62
Branch Instruction BI Field .....	63
Branch Instruction BO Field .....	63
Branch Prediction .....	64
Branch Control Registers .....	65
Link Register (LR) .....	65
Count Register (CTR) .....	66
Condition Register (CR) .....	66
Integer Processing .....	69
General Purpose Registers (GPRs) .....	69
Integer Exception Register (XER) .....	70
Summary Overflow (SO) Field .....	71
Overflow (OV) Field .....	71
Carry (CA) Field .....	72
Processor Control .....	72
Special Purpose Registers General (USPRG0, SPRG0–SPRG7) .....	73
Processor Version Register (PVR) .....	73
Processor Identification Register (PIR) .....	74
Core Configuration Register 0 (CCR0) .....	74
Core Configuration Register 1 (CCR1) .....	76
Reset Configuration (RSTCFG) .....	77
User and Supervisor Modes .....	78
Privileged Instructions .....	78
Privileged SPRs .....	79
Speculative Accesses .....	79
Synchronization .....	79
Context Synchronization .....	80

Execution Synchronization .....	81
Storage Ordering and Synchronization .....	81
<b>3. Initialization .....</b>	<b>83</b>
PPC440x5 Core State After Reset .....	83
Reset Types .....	87
Reset Sources .....	87
Initialization Software Requirements .....	87
<b>4. Instruction and Data Caches .....</b>	<b>93</b>
Cache Array Organization and Operation .....	93
Cache Line Replacement Policy .....	94
Cache Locking and Transient Mechanism .....	96
Instruction Cache Controller .....	100
ICC Operations .....	101
Speculative Prefetch Mechanism .....	102
Instruction Cache Coherency .....	103
Self-Modifying Code .....	103
Instruction Cache Synonyms .....	104
Instruction Cache Control and Debug .....	105
Instruction Cache Management and Debug Instruction Summary .....	105
Core Configuration Register 0 (CCR0) .....	105
Core Configuration Register 1 (CCR1) .....	107
icbt Operation .....	108
icread Operation .....	109
Instruction Cache Parity Operations .....	110
Simulating Instruction Cache Parity Errors for Software Testing .....	111
Data Cache Controller .....	112
DCC Operations .....	113
Load and Store Alignment .....	114
Load Operations .....	115
Store Operations .....	115
Line Flush Operations .....	117
Data Read PLB Interface Requests .....	118
Data Write PLB Interface Requests .....	119
Storage Access Ordering .....	120
Data Cache Coherency .....	120
Data Cache Control and Debug .....	121
Data Cache Management and Debug Instruction Summary .....	121
Core Configuration Register 0 (CCR0) .....	122
Core Configuration Register 1 (CCR1) .....	122
dcbt and dcbtst Operation .....	122
dcread Operation .....	123
Data Cache Parity Operations .....	125
Simulating Data Cache Parity Errors for Software Testing .....	126
<b>5. Memory Management .....</b>	<b>129</b>
MMU Overview .....	129
Support for PowerPC Book-E MMU Architecture .....	129
Translation Lookaside Buffer .....	130

Page Identification .....	134
Virtual Address Formation .....	134
Address Space Identifier Convention .....	134
TLB Match Process .....	135
Address Translation .....	136
Access Control .....	138
Execute Access .....	138
Write Access .....	138
Read Access .....	139
Access Control Applied to Cache Management Instructions .....	139
Storage Attributes .....	140
Write-Through (W) .....	141
Caching Inhibited (I) .....	141
Memory Coherence Required (M) .....	141
Guarded (G) .....	141
Endian (E) .....	142
User-Definable (U0–U3) .....	142
Supported Storage Attribute Combinations .....	143
Storage Control Registers .....	143
Memory Management Unit Control Register (MMUCR) .....	143
Process ID (PID) .....	146
Shadow TLB Arrays .....	147
TLB Management Instructions .....	148
TLB Search Instruction (tlbsx[.]) .....	148
TLB Read/Write Instructions (tlbre/tlbwe) .....	149
TLB Sync Instruction (tlbsync) .....	149
Page Reference and Change Status Management .....	150
TLB Parity Operations .....	150
Reading TLB Parity Bits with tlbre .....	151
Simulating TLB Parity Errors for Software Testing .....	152
<b>6. Interrupts and Exceptions .....</b>	<b>153</b>
Overview .....	153
Interrupt Classes .....	153
Asynchronous Interrupts .....	153
Synchronous Interrupts .....	153
Synchronous, Precise Interrupts .....	154
Synchronous, Imprecise Interrupts .....	154
Critical and Non-Critical Interrupts .....	155
Machine Check Interrupts .....	155
Interrupt Processing .....	156
Partially Executed Instructions .....	158
Interrupt Processing Registers .....	159
Machine State Register (MSR) .....	159
Save/Restore Register 0 (SRR0) .....	161
Save/Restore Register 1 (SRR1) .....	161
Critical Save/Restore Register 0 (CSRR0) .....	162
Critical Save/Restore Register 1 (CSRR1) .....	162
Machine Check Save/Restore Register 0 (MCSRR0) .....	163
Machine Check Save/Restore Register 1 (MCSRR1) .....	163

Data Exception Address Register (DEAR) .....	164
Interrupt Vector Offset Registers (IVOR0–IVOR15) .....	164
Interrupt Vector Prefix Register (IVPR) .....	165
Exception Syndrome Register (ESR) .....	166
Machine Check Status Register (MCSR) .....	168
Interrupt Definitions .....	169
Critical Input Interrupt .....	172
Machine Check Interrupt .....	172
Data Storage Interrupt .....	175
Instruction Storage Interrupt .....	178
External Input Interrupt .....	179
Alignment Interrupt .....	179
Program Interrupt .....	180
Floating-Point Unavailable Interrupt .....	183
System Call Interrupt .....	184
Auxiliary Processor Unavailable Interrupt .....	184
Decrementer Interrupt .....	185
Fixed-Interval Timer Interrupt .....	185
Watchdog Timer Interrupt .....	186
Data TLB Error Interrupt .....	186
Instruction TLB Error Interrupt .....	188
Debug Interrupt .....	188
Interrupt Ordering and Masking .....	192
Interrupt Ordering Software Requirements .....	193
Interrupt Order .....	194
Exception Priorities .....	195
Exception Priorities for Integer Load, Store, and Cache Management Instructions .....	196
Exception Priorities for Floating-Point Load and Store Instructions .....	196
Exception Priorities for Allocated Load and Store Instructions .....	197
Exception Priorities for Floating-Point Instructions (Other) .....	197
Exception Priorities for Allocated Instructions (Other) .....	198
Exception Priorities for Privileged Instructions .....	199
Exception Priorities for Trap Instructions .....	199
Exception Priorities for System Call Instruction .....	200
Exception Priorities for Branch Instructions .....	200
Exception Priorities for Return From Interrupt Instructions .....	200
Exception Priorities for Preserved Instructions .....	200
Exception Priorities for Reserved Instructions .....	201
Exception Priorities for All Other Instructions .....	201
<b>7. Timer Facilities .....</b>	<b>203</b>
Time Base .....	204
Reading the Time Base .....	204
Writing the Time Base .....	205
Decrementer (DEC) .....	205
Fixed Interval Timer (FIT) .....	206
Watchdog Timer .....	207
Timer Control Register (TCR) .....	209
Timer Status Register (TSR) .....	210
Freezing the Timer Facilities .....	211

Selection of the Timer Clock Source .....	211
<b>8. Debug Facilities .....</b>	<b>213</b>
Support for Development Tools .....	213
Debug Modes .....	213
Internal Debug Mode .....	213
External Debug Mode .....	214
Debug Wait Mode .....	214
Trace Debug Mode .....	215
Debug Events .....	215
Instruction Address Compare (IAC) Debug Event .....	216
IAC Debug Event Fields .....	216
IAC Debug Event Processing .....	219
Data Address Compare (DAC) Debug Event .....	220
DAC Debug Event Fields .....	220
DAC Debug Event Processing .....	223
DAC Debug Events Applied to Instructions that Result in Multiple Storage Accesses .....	224
DAC Debug Events Applied to Various Instruction Types .....	224
Data Value Compare (DVC) Debug Event .....	225
DVC Debug Event Fields .....	226
DVC Debug Event Processing .....	227
DVC Debug Events Applied to Instructions that Result in Multiple Storage Accesses .....	227
DVC Debug Events Applied to Various Instruction Types .....	227
Branch Taken (BRT) Debug Event .....	228
Trap (TRAP) Debug Event .....	228
Return (RET) Debug Event .....	229
Instruction Complete (ICMP) Debug Event .....	229
Interrupt (IRPT) Debug Event .....	230
Unconditional Debug Event (UDE) .....	231
Debug Event Summary .....	231
Debug Reset .....	232
Debug Timer Freeze .....	232
Debug Registers .....	232
Debug Control Register 0 (DBCR0) .....	233
Debug Control Register 1 (DBCR1) .....	234
Debug Control Register 2 (DBCR2) .....	237
Debug Status Register (DBSR) .....	238
Instruction Address Compare Registers (IAC1–IAC4) .....	239
Data Address Compare Registers (DAC1–DAC2) .....	240
Data Value Compare Registers (DVC1–DVC2) .....	240
Debug Data Register (DBDR) .....	241
<b>9. Instruction Set .....</b>	<b>243</b>
Instruction Set Portability .....	244
Instruction Formats .....	244
Pseudocode .....	245
Operator Precedence .....	247
Register Usage .....	247
Alphabetical Instruction Listing .....	247
add .....	248





---

addc .....	249
adde .....	250
addi .....	251
addic .....	252
addic. ....	253
addis .....	254
addme .....	255
addze .....	256
and .....	257
andc .....	258
andi. ....	259
andis. ....	260
b .....	261
bc .....	262
bcctr .....	268
bclr .....	271
cmp .....	275
cmpi .....	276
cmpl .....	277
cmpli .....	278
cntlzw .....	279
crand .....	280
crandc .....	281
creqv .....	282
crnand .....	283
crnor .....	284
cror .....	285
crorc .....	286
crxor .....	287
dcba .....	288
dcbf .....	289
dcbi .....	290
dcbst .....	291
dcbt .....	292
dcbtst .....	293
dcbz .....	294
dccci .....	295
dcread .....	296
divw .....	298
divwu .....	299
dlmzb .....	300
eqv .....	301
extsb .....	302
extsh .....	303
icbi .....	304
icbt .....	305
iccci .....	307
icread .....	308
isel .....	310
isync .....	311
lbz .....	312
lbzu .....	313



lbzux .....	314
lbzx .....	315
lha .....	316
lhau .....	317
lhaux .....	318
lhax .....	319
lhbrx .....	320
lhz .....	321
lhzu .....	322
lhzux .....	323
lhzx .....	324
lmw .....	325
lswi .....	326
lswx .....	328
lwarx .....	330
lwbrx .....	331
lwz .....	332
lwzu .....	333
lwzux .....	334
lwzx .....	335
macchw .....	336
macchws .....	337
macchwsu .....	338
macchwu .....	339
machhw .....	340
machhws .....	341
machhwsu .....	342
machhwu .....	343
macihw .....	344
macihws .....	345
macihwsu .....	346
macihwu .....	347
mbar .....	348
mcrf .....	350
mcrxr .....	351
mfcrl .....	352
mfdcr .....	353
mfmsr .....	354
mfspr .....	355
msync .....	358
mtcrf .....	359
mtdcr .....	360
mtmsr .....	361
mtspr .....	362
mulchw .....	365
mulchwu .....	366
mulhhw .....	367
mulhhwu .....	368
mulhw .....	369
mulhwu .....	370
mullhw .....	371
mullhwu .....	372



---

mulli .....	373
mullw .....	374
nand .....	375
neg .....	376
nmacchw .....	377
nmacchws .....	378
nmachhw .....	379
nmachhws .....	380
nmaclhw .....	381
nmaclhws .....	382
nor .....	383
or .....	384
orc .....	385
ori .....	386
oris .....	387
rfci .....	388
rfi .....	389
rfmci .....	390
rlwimi .....	391
rlwinm .....	392
rlwnm .....	394
sc .....	395
slw .....	396
sraw .....	397
srawi .....	398
srw .....	399
stb .....	400
stbu .....	401
stbux .....	402
stbx .....	403
sth .....	404
sthbrx .....	405
sthu .....	406
sthux .....	407
sthx .....	408
stmw .....	409
stswi .....	410
stswx .....	411
stw .....	412
stwbrx .....	413
stwcx .....	414
stwu .....	416
stwux .....	417
stwx .....	418
subf .....	419
subfc .....	420
subfe .....	421
subfic .....	422
subfme .....	423
subfze .....	424
tlbre .....	425
tlbsx .....	427



tlbsync .....	428
tlbwe .....	429
tw .....	430
twi .....	433
wrtee .....	435
wrteei .....	436
xor .....	437
xori .....	438
xoris .....	439
<b>10. Register Summary .....</b>	<b>441</b>
Register Categories .....	441
Reserved Fields .....	445
Device Control Registers .....	445
Alphabetical Register Listing .....	447
CCR0 .....	448
CCR1 .....	450
CR .....	452
CSRR0 .....	453
CSRR1 .....	454
CTR .....	455
DAC1–DAC2 .....	456
DBCR0 .....	457
DBCR1 .....	459
DBCR2 .....	461
DBDR .....	463
DBSR .....	464
DCDBTRH .....	466
DCDBTRL .....	467
DEAR .....	468
DEC .....	469
DECAR .....	470
DNV0–DNV3 .....	471
DTV0–DTV3 .....	472
DVC1–DVC2 .....	473
DVLIM .....	474
ESR .....	475
GPR0–GPR31 .....	477
IAC1–IAC4 .....	478
ICDBDR .....	479
ICDBTRH .....	480
ICDBTRL .....	481
INV0–INV3 .....	482
ITV0–ITV3 .....	483
IVLIM .....	484
IVOR0–IVOR15 .....	485
IVPR .....	486
LR .....	487
MCSR .....	488
MCSRR0 .....	489
MCSRR1 .....	490



---

MMUCR.....	491
MSR.....	492
PID.....	494
PIR.....	495
PVR.....	496
RSTCFG.....	497
SPRG0–SPRG7.....	498
SRR0.....	499
SRR1.....	500
TBL.....	501
TBU.....	502
TCR.....	503
TSR.....	504
USPRG0.....	505
XER.....	506
<b>Appendix A. Instruction Summary .....</b>	<b>507</b>
Instruction Formats .....	507
Instruction Fields .....	508
Instruction Format Diagrams .....	509
I-Form .....	510
B-Form .....	510
SC-Form .....	510
D-Form .....	510
X-Form .....	511
XL-Form .....	512
XFX-Form .....	512
XO-Form .....	512
M-Form .....	512
Alphabetical Summary of Implemented Instructions .....	512
Allocated Instruction Opcodes .....	543
Preserved Instruction Opcodes .....	543
Reserved Instruction Opcodes .....	544
Implemented Instructions Sorted by Opcode .....	544
<b>Appendix B. PPC440x5 Core Compiler Optimizations .....</b>	<b>553</b>
<b>Index .....</b>	<b>555</b>
<b>Revision Log .....</b>	<b>573</b>



## Figures

Figure 1-1.	PPC440x5 Core Block Diagram .....	30
Figure 2-1.	User Programming Model Registers .....	47
Figure 2-2.	Supervisor Programming Model Registers .....	49
Figure 2-3.	Link Register (LR) .....	65
Figure 2-4.	Count Register (CTR) .....	66
Figure 2-5.	Condition Register (CR) .....	66
Figure 2-6.	General Purpose Registers (R0-R31) .....	69
Figure 2-7.	Integer Exception Register (XER) .....	70
Figure 2-8.	Special Purpose Registers General (USPRG0, SPRG0–SPRG7) .....	73
Figure 2-9.	Processor Version Register (PVR) .....	74
Figure 2-10.	Processor Identification Register (PIR) .....	74
Figure 2-11.	Core Configuration Register 0 (CCR0) .....	75
Figure 2-12.	Core Configuration Register 1 (CCR1) .....	76
Figure 2-13.	Reset Configuration .....	77
Figure 4-1.	Instruction Cache Normal Victim Registers (INV0–INV3) .....	95
Figure 4-1.	Instruction Cache Transient Victim Registers (ITV0–ITV3) .....	95
Figure 4-1.	Data Cache Normal Victim Registers (DNV0–DNV3) .....	95
Figure 4-1.	Data Cache Transient Victim Registers (DTV0–DTV3) .....	95
Figure 4-2.	Instruction Cache Victim Limit (IVLIM) .....	97
Figure 4-2.	Data Cache Victim Limit (DVLIM) .....	97
Figure 4-3.	Cache Locking and Transient Mechanism (Example 1)1 .....	99
Figure 4-4.	Cache Locking and Transient Mechanism (Example 2) .....	100
Figure 4-5.	Core Configuration Register 0 (CCR0) .....	106
Figure 4-6.	Core Configuration Register 1 (CCR1) .....	107
Figure 4-7.	Instruction Cache Debug Data Register (ICDBDR) .....	109
Figure 4-8.	Instruction Cache Debug Tag Register High (ICDBTRH) .....	110
Figure 4-9.	Instruction Cache Debug Tag Register Low (ICDBTRL) .....	110
Figure 4-10.	Data Cache Debug Tag Register High (DCDBTRH) .....	124
Figure 4-11.	Data Cache Debug Tag Register Low (DCDBTRL) .....	124
Figure 5-1.	Virtual Address to TLB Entry Match Process .....	136
Figure 5-2.	Effective-to-Real Address Translation Flow .....	137
Figure 5-3.	Memory Management Unit Control Register (MMUCR) .....	143
Figure 5-4.	Process ID (PID) .....	147
Figure 5-5.	TLB Entry Word Definitions .....	149
Figure 6-1.	Machine State Register (MSR) .....	159
Figure 6-2.	Save/Restore Register 0 (SRR0) .....	161
Figure 6-3.	Save/Restore Register 1 (SRR1) .....	162
Figure 6-4.	Critical Save/Restore Register 0 (CSRR0) .....	162

Figure 6-5.	Critical Save/Restore Register 1 (CSRR1) .....	163
Figure 6-6.	Machine Check Save/Restore Register 0 (MCSRR0) .....	163
Figure 6-7.	Machine Check Save/Restore Register 1 (MCSRR1) .....	164
Figure 6-8.	Data Exception Address Register (DEAR) .....	164
Figure 6-9.	Interrupt Vector Offset Registers (IVOR0–IVOR15) .....	165
Figure 6-10.	Interrupt Vector Prefix Register (IVPR) .....	166
Figure 6-11.	Exception Syndrome Register (ESR) .....	166
Figure 6-12.	Machine Check Status Register (MCSR) .....	168
Figure 7-1.	Relationship of Timer Facilities to the Time Base .....	203
Figure 7-2.	Time Base Lower (TBL) .....	204
Figure 7-3.	Time Base Upper (TBU) .....	204
Figure 7-4.	Decrementer (DEC) .....	206
Figure 7-5.	Decrementer Auto-Reload (DECAR) .....	206
Figure 7-6.	Watchdog State Machine .....	209
Figure 7-7.	Timer Control Register (TCR) .....	210
Figure 7-8.	Timer Status Register (TSR) .....	211
Figure 8-1.	Debug Control Register 0 (DBCR0) .....	233
Figure 8-2.	Debug Control Register 1 (DBCR1) .....	234
Figure 8-3.	Debug Control Register 2 (DBCR2) .....	237
Figure 8-4.	Debug Status Register (DBSR) .....	238
Figure 8-5.	Instruction Address Compare Registers (IAC1–IAC4) .....	240
Figure 8-6.	Data Address Compare Registers (DAC1–DAC2) .....	240
Figure 8-7.	Data Value Compare Registers (DVC1–DVC2) .....	240
Figure 8-8.	Debug Data Register (DBDR) .....	241
Figure 10-1.	Core Configuration Register 0 (CCR0) .....	448
Figure 10-2.	Core Configuration Register 1 (CCR1) .....	450
Figure 10-3.	Condition Register (CR) .....	452
Figure 10-4.	Critical Save/Restore Register 0 (CSRR0) .....	453
Figure 10-5.	Critical Save/Restore Register 1 (CSRR1) .....	454
Figure 10-6.	Count Register (CTR) .....	455
Figure 10-7.	Data Address Compare Registers (DAC1–DAC2) .....	456
Figure 10-8.	Debug Control Register 0 (DBCR0) .....	457
Figure 10-9.	Debug Control Register 1 (DBCR1) .....	459
Figure 10-10.	Debug Control Register 2 (DBCR2) .....	461
Figure 10-11.	Debug Data Register (DBDR) .....	463
Figure 10-12.	Debug Status Register (DBSR) .....	464
Figure 10-13.	Data Cache Debug Tag Register High (DCDBTRH) .....	466
Figure 10-14.	Data Cache Debug Tag Register Low (DCDBTRL) .....	467
Figure 10-15.	Data Exception Address Register (DEAR) .....	468



Figure 10-16. Decrementer (DEC) .....	469
Figure 10-17. Decrementer Auto-Reload (DECAR) .....	470
Figure 10-18. Data Cache Normal Victim Registers (DNV0–DNV3) .....	471
Figure 10-19. Data Cache Transient Victim Registers (DTV0–DTV3) .....	472
Figure 10-20. Data Value Compare Registers (DVC1–DVC2) .....	473
Figure 10-21. Data Cache Victim Limit (DVLIM) .....	474
Figure 10-22. Exception Syndrome Register (ESR) .....	475
Figure 10-23. General Purpose Registers (R0-R31) .....	477
Figure 10-24. Instruction Address Compare Registers (IAC1–IAC4) .....	478
Figure 10-25. Instruction Cache Debug Data Register (ICDBDR) .....	479
Figure 10-26. Instruction Cache Debug Tag Register High (ICDBTRH) .....	480
Figure 10-27. Instruction Cache Debug Tag Register Low (ICDBTRL) .....	481
Figure 10-28. Instruction Cache Normal Victim Registers (INV0–INV3) .....	482
Figure 10-29. Instruction Cache Transient Victim Registers (ITV0–ITV3) .....	483
Figure 10-30. Instruction Cache Victim Limit (IVLIM) .....	484
Figure 10-31. Interrupt Vector Offset Registers (IVOR0–IVOR15) .....	485
Figure 10-32. Interrupt Vector Prefix Register (IVPR) .....	486
Figure 10-33. Link Register (LR) .....	487
Figure 10-34. Machine Check Status Register (MCSR) .....	488
Figure 10-35. Machine Check Save/Restore Register 0 (MCSRR0) .....	489
Figure 10-36. Machine Check Save/Restore Register 1 (MCSRR1) .....	490
Figure 10-37. Memory Management Unit Control Register (MMUCR) .....	491
Figure 10-38. Machine State Register (MSR) .....	492
Figure 10-39. Process ID (PID) .....	494
Figure 10-40. Processor Identification Register (PIR) .....	495
Figure 10-41. Processor Version Register (PVR) .....	496
Figure 10-42. Reset Configuration .....	497
Figure 10-43. Special Purpose Registers General (SPRG0–SPRG7) .....	498
Figure 10-44. Save/Restore Register 0 (SRR0) .....	499
Figure 10-45. Save/Restore Register 1 (SRR1) .....	500
Figure 10-46. Time Base Lower (TBL) .....	501
Figure 10-47. Time Base Upper (TBU) .....	502
Figure 10-48. Timer Control Register (TCR) .....	503
Figure 10-49. Timer Status Register (TSR) .....	504
Figure 10-50. User Special Purpose Register General (USPRG0) .....	505
Figure 10-51. Integer Exception Register (XER) .....	506
Figure A-1. I Instruction Format .....	510
Figure A-2. B Instruction Format .....	510
Figure A-3. SC Instruction Format .....	510



**PPC440x5 CPU Core**

**Preliminary**

---

Figure A-4.	D Instruction Format .....	510
Figure A-5.	X Instruction Format .....	511
Figure A-6.	XL Instruction Format .....	512
Figure A-7.	XFX Instruction Format .....	512
Figure A-8.	XO Instruction Format .....	512
Figure A-9.	M Instruction Format .....	512

## Tables

Table 2-1.	Data Operand Definitions .....	40
Table 2-2.	Alignment Effects for Storage Access Instructions .....	40
Table 2-3.	Register Categories .....	50
Table 2-4.	Instruction Categories .....	56
Table 2-5.	Integer Storage Access Instructions .....	57
Table 2-6.	Integer Arithmetic Instructions .....	57
Table 2-7.	Integer Logical Instructions .....	58
Table 2-8.	Integer Compare Instructions .....	58
Table 2-9.	Integer Trap Instructions .....	58
Table 2-10.	Integer Rotate Instructions .....	58
Table 2-11.	Integer Shift Instructions .....	59
Table 2-12.	Integer Select Instruction .....	59
Table 2-13.	Branch Instructions .....	59
Table 2-14.	Condition Register Logical Instructions .....	60
Table 2-15.	Register Management Instructions .....	60
Table 2-16.	System Linkage Instructions .....	60
Table 2-17.	Processor Synchronization Instruction .....	60
Table 2-18.	Cache Management Instructions .....	61
Table 2-19.	TLB Management Instructions .....	61
Table 2-20.	Storage Synchronization Instructions .....	62
Table 2-21.	Allocated Instructions .....	62
Table 2-22.	BO Field Definition .....	63
Table 2-23.	BO Field Examples .....	64
Table 2-24.	CR Updating Instructions .....	67
Table 2-25.	XER[SO,OV] Updating Instructions .....	71
Table 2-26.	XER[CA] Updating Instructions .....	71
Table 2-27.	Privileged Instructions .....	78
Table 3-1.	Reset Values of Registers and Other PPC440x5 Facilities .....	84
Table 4-1.	Instruction and Data Cache Array Organization .....	94
Table 4-2.	Cache Sizes and Parameters .....	94
Table 4-3.	Victim Index Field Selection .....	96
Table 4-4.	Icread and dcread Cache Line Selection .....	109
Table 4-5.	Data Cache Behavior on Store Accesses .....	117
Table 5-1.	TLB Entry Fields .....	131
Table 5-2.	Page Size and Effective Address to EPN Comparison .....	136
Table 5-3.	Page Size and Real Address Formation .....	137
Table 5-4.	Access Control Applied to Cache Management Instructions .....	140
Table 6-1.	Interrupt Types Associated with each IVOR .....	165

Table 6-2.	Interrupt and Exception Types .....	169
Table 7-1.	Fixed Interval Timer Period Selection .....	207
Table 7-2.	Watchdog Timer Period Selection .....	207
Table 7-3.	Watchdog Timer Exception Behavior .....	208
Table 8-1.	Debug Events .....	215
Table 8-2.	IAC Range Mode Auto-Toggle Summary .....	219
Table 8-3.	Debug Event Summary .....	231
Table 9-1.	Instruction Categories .....	243
Table 9-2.	Allocated Instructions .....	244
Table 9-3.	Operator Precedence .....	247
Table 9-4.	Extended Mnemonics for addi .....	251
Table 9-5.	Extended Mnemonics for addic .....	252
Table 9-6.	Extended Mnemonics for addic. ....	253
Table 9-7.	Extended Mnemonics for addis .....	254
Table 9-8.	Extended Mnemonics for bc, bca, bcl, bcla .....	263
Table 9-9.	Extended Mnemonics for bcctr, bcctrl .....	268
Table 9-10.	Extended Mnemonics for bclr, bclrl .....	272
Table 9-11.	Extended Mnemonics for cmp .....	275
Table 9-12.	Extended Mnemonics for cmpi .....	276
Table 9-13.	Extended Mnemonics for cmpl .....	277
Table 9-14.	Extended Mnemonics for cmpli .....	278
Table 9-15.	Extended Mnemonics for creqv .....	282
Table 9-16.	Extended Mnemonics for crnor .....	284
Table 9-17.	Extended Mnemonics for cror .....	285
Table 9-18.	Extended Mnemonics for crxor .....	287
Table 9-19.	Extended Mnemonics for mbar .....	348
Table 9-20.	Extended Mnemonics for mfspr .....	356
Table 9-21.	FXM Bit Field Correspondence .....	359
Table 9-22.	Extended Mnemonics for mtrf .....	359
Table 9-23.	Extended Mnemonics for mtspr .....	363
Table 9-24.	Extended Mnemonics for nor, nor. ....	383
Table 9-25.	Extended Mnemonics for or, or. ....	384
Table 9-26.	Extended Mnemonics for ori .....	386
Table 9-27.	Extended Mnemonics for rlwimi, rlwimi. ....	391
Table 9-28.	Extended Mnemonics for rlwinm, rlwinm. ....	392
Table 9-29.	Extended Mnemonics for rlwnm, rlwnm. ....	394
Table 9-30.	Extended Mnemonics for subf, subf., subfo, subfo. ....	419
Table 9-31.	Extended Mnemonics for subfc, subfc., subfco, subfco. ....	420
Table 9-32.	Extended Mnemonics for tw .....	431



---

Table 9-33.	Extended Mnemonics for twi .....	434
Table 10-1.	Register Categories .....	441
Table 10-2.	Special Purpose Registers Sorted by SPR Number .....	443
Table 10-3.	Interrupt Types Associated with each IVOR .....	485
Table A-1.	PPC440x5 Instruction Syntax Summary .....	513
Table A-2.	Allocated Opcodes .....	543
Table A-3.	Preserved Opcodes .....	543
Table A-4.	Reserved-nop Opcodes .....	544
Table A-5.	PPC440x5 Instructions by Opcode .....	545



## About This Book

This user's manual provides the architectural overview, programming model, and detailed information about the instruction set, registers, and other facilities of the IBM™ Book-E Enhanced PowerPC™ 440x5 (PPC440x5™) 32-bit embedded controller core.

The PPC440x5 embedded controller core features:

- Book-E Enhanced PowerPC Architecture™
- Dual-issue superscalar pipeline with dynamic branch prediction
- Separate, configurable (up to 32KB each) instruction and data caches, with cache line locking
- DSP acceleration with 24 new integer multiply-accumulate (MAC) instructions
- Memory Management Unit (MMU) with 64-entry TLB and support for page sizes of 1KB–256MB
- 64GB (36-bit) physical address capability
- 128-bit PLB interface, part of the IBM CoreConnect™ on-chip system bus architecture
- JTAG debug interface with extensive integrated debug facilities, including real-time trace

## Who Should Use This Book

This book is for system hardware and software developers, and for application developers who need to understand the PPC440x5. The audience should understand embedded system design, operating systems, RISC microprocessing, and computer organization and architecture.

---

## How to Use This Book

This book describes the PPC440x5 device architecture, programming model, registers, and instruction set. This book contains the following chapters:

- *Overview* on page 27
- *Programming Model* on page 39
- *Initialization* on page 83
- *Instruction and Data Caches* on page 93
- *Memory Management* on page 129
- *Interrupts and Exceptions* on page 153
- *Timer Facilities* on page 203
- *Debug Facilities* on page 213
- *Instruction Set* on page 243
- *Register Summary* on page 441

This book contains the following appendixes:

- *Instruction Summary* on page 507
- *PPC440x5 Core Compiler Optimizations* on page 553

To help readers find material in these chapters, this book contains:

- *Contents* on page 3.
- *Figures* on page 15.
- *Tables* on page 19.
- *Index* on page 555.



## Notation

The manual uses the following notational conventions:

- Active low signals are shown with an overbar ( $\overline{\text{Active\_Low}}$ )
- All numbers are decimal unless specified in some special way.
  - 0bnnnn means a number expressed in binary format.
  - 0xnxxx means a number expressed in hexadecimal format.

Underscores may be used between digits.

- RA refers to General Purpose Register (GPR) RA.
- (RA) refers to the contents of GPR RA.
- (RA|0) refers to the contents of GPR RA, or to the value 0 if the RA field is 0.
- Bits in registers, instructions, and fields are specified as follows.
  - Bits are numbered most-significant bit to least-significant bit, starting with bit 0.

**Note:** This document differs from the Book-E architecture specification in the use of bit numbering for architected registers. Book-E defines the full, 64-bit instruction set architecture, and all registers are shown as having bit numbers from 0 to 63, with bit 63 being the least significant. This manual describes a 32-bit subset implementation of the architecture. Architected registers are described as being 32 bits long, with bits numbered from 0 to 31, and with bit 31 being the least significant. When this document refers to register bits 0 to 31, they actually correspond to bits 32 to 63 of the same register in the Book-E architecture specification.

- $X_p$  means bit p of register, instruction, or field X
- $X_{p:q}$  means bits p through q of register, instruction, or field X
- $X_{p,q,\dots}$  means bits p, q, ... of register, instruction, or field X
- $X[p]$  means a named field p of register X.
- $X[p:q]$  means named fields p through q of register X.
- $X[p,q,\dots] \dots$  means named fields p, q, ... of register X.
- $\neg X$  means the ones complement of the contents of X.
- A period (.) as the last character of an instruction mnemonic means that the instruction records status information in certain fields of the Condition Register as a side effect of execution, as described in Chapter 9, "Instruction Set."
- The symbol || is used to describe the concatenation of two values. For example, 0b010 || 0b111 is the same as 0b010111.
- $x^n$  means x raised to the n power.
- ${}^n x$  means the replication of x, n times (that is, 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 0b00000.
  - ${}^n 1$  means a field of n bits with each bit equal to 1. Thus  ${}^5 1$  is equivalent to 0b11111.
- /, //, ///, ... denotes a reserved field in an instruction or in a register.
- ? denotes an allocated bit in a register.

- A shaded field denotes a field that is reserved or allocated in an instruction or in a register.

## Related Publications

The following book describes the Book-E Enhanced PowerPC Architecture:

- Book E: PowerPC Architecture Enhanced for Embedded Applications ([www.chips.ibm.com/techlib/products/powerpc/manuals/](http://www.chips.ibm.com/techlib/products/powerpc/manuals/))

The following CD-ROM contains publications describing the IBM PowerPC 400 family of embedded controllers, including this manual *PowerPC PPC440x5 User's Manual*, and application and technical notes.

- IBM PowerPC Embedded Processor Solutions (Order Number SC09-3032)

## 1. Overview

The IBM™ PowerPC™ 440x5 32-bit embedded processor core, referred to as the PPC440x5 core, implements the Book-E Enhanced PowerPC Architecture.

This chapter describes:

- PPC440x5 core features
- The PPC440x5 core as an implementation of the Book-E Enhanced PowerPC Architecture
- The organization of the PPC440x5 core, including a block diagram and descriptions of the functional units
- PPC440x5 core interfaces

### 1.1 PPC440x5 Features

The PPC440x5 core is a high-performance, low-power engine that implements the flexible and powerful Book-E Enhanced PowerPC Architecture.

The PPC440x5 contains a dual-issue, superscalar, pipelined processing unit, along with other functional elements required by embedded ASIC product specifications. These other functions include memory management, cache control, timers, and debug facilities. Interfaces for custom co-processors and floating point functions are provided, along with separate instruction and data cache array interfaces which can be configured to various sizes (optimized for 32KB). The processor local bus (PLB) system interface has been extended to 128 bits and is fully compatible with the IBM CoreConnect on-chip system architecture, providing the framework to efficiently support system-on-a-chip (SOC) designs.

In addition, the PPC440x5 core is a member of the PowerPC 400 Series of advanced embedded processors cores, which is supported by the PowerPC Embedded Tools Program. In this program, IBM and many third-party vendors offer a full range of robust development tools for embedded applications. Among these are compilers, debuggers, real-time operating systems, and logic analyzers.

PPC440x5 features include:

- High performance, dual-issue, superscalar 32-bit RISC CPU
  - Superscalar implementation of the full 32-bit Book-E Enhanced PowerPC Architecture
  - Seven stage, highly-pipelined micro-architecture
  - Dual instruction fetch, decode, and out-of-order issue
  - Out-of-order dispatch, execution, and completion
  - High-accuracy dynamic branch prediction using a Branch History Table (BHT)
  - Reduced branch latency using Branch Target Address Cache (BTAC)
  - Three independent pipelines
    - Combined complex integer, system, and branch pipeline
    - Simple integer pipeline
    - Load/store pipeline
  - Single cycle multiply
  - Single cycle multiply-accumulate (DSP instruction set extensions)

- 9-port (6-read, 3-write) 32x32-bit General Purpose Register (GPR) file
- Hardware support for all CPU misaligned accesses
- Full support for both big and little endian byte ordering
- Extensive power management designed into core for maximum performance/power efficiency
- Primary caches
  - Independently configurable instruction and data cache arrays
  - Array size offerings: 32KB, 16KB, and 8KB
  - Single-cycle access
  - 32-byte (eight word) line size
  - Highly-associative (64-way for 32KB/16KB, 32-way for 8KB)
  - Write-back and write-through operation
  - Control over whether stores will allocate or write-through on cache miss
  - Extensive load/store queues and multiple line fill/flush buffers
  - Non-blocking with up to four outstanding load misses
  - Cache line locking supported
  - Caches can be partitioned to provide separate regions for “transient” instructions and data
    - High associativity permits efficient allocation of cache memory
  - Critical word first data access and forwarding
  - Cache tags and data are parity protected against soft errors.
- Memory Management Unit
  - Separate instruction and data shadow TLBs
  - 64-entry, fully-associative unified TLB array
  - Variable page sizes (1KB-256MB), simultaneously resident in TLB
  - 4-bit extended real address for 36-bit (64 GB) addressability
  - Flexible TLB management with software page table search
  - Storage attribute controls for write-through, caching inhibited, guarded, and byte order (endianness)
  - Four user-definable storage attribute controls (for controlling CodePack™ code compression and transient data, for example)
  - TLB tags and data are parity protected against soft errors.
- Debug facilities
  - Extensive hardware debug facilities incorporated into the IEEE 1149.1 JTAG port
    - Multiple instruction and data address breakpoints (including range)
    - Data value compare
    - Single-step, branch, trap, and other debug events
  - Non-invasive real-time software trace interface
- Timer facilities
  - 64-bit time base

- Decrementer with auto-reload capability
- Fixed Interval Timer (FIT)
- Watchdog Timer with critical interrupt and/or auto-reset
- Multiple core Interfaces defined by the IBM CoreConnect on-chip system architecture
  - PLB interfaces
    - Three independent 128-bit interfaces for instruction reads, data reads, and data writes
    - Glueless attachment to 32-, 64-, or 128-bit CoreConnect system environments
    - Multiple CPU:PLB frequency ratios supported (N:1, N:2, N:3)
    - 6.4 GB/sec maximum data rate to CPU
    - On-chip memory (OCM) integration capability over the PLB interface
  - Auxiliary Processor Unit (APU) Port
    - Provides functional extensions to the processor pipelines, including GPR file operations
    - 128-bit load/store interface (direct access between APU and the primary data cache)
    - Interface can support APU execution of all PowerPC floating point instructions
    - Attachment capability for DSP co-processing such as accumulators and SIMD computation
    - Enables customer-specific instruction enhancements for multimedia applications
  - Device Control Register (DCR) interface for independent access to on-chip control registers
    - Avoids contention for high-bandwidth PLB system bus
  - Clock and power management interface
  - JTAG debug interface

## 1.2 The PPC440x5 as a PowerPC Implementation

The PPC440x5 core implements the full, 32-bit fixed-point subset of the Book-E Enhanced PowerPC Architecture. The PPC440x5 core fully complies with these architectural specifications. The 64-bit operations of the architecture are not supported, and the core does not implement the floating point operations, although a floating point unit (FPU) may be attached (using the APU interface). Within the core, the 64-bit operations and the floating point operations are trapped, and the floating point operations can be emulated using software.

See Appendix A of the Book-E Enhanced PowerPC Architecture specification for more information on 32-bit subset implementations of the architecture.

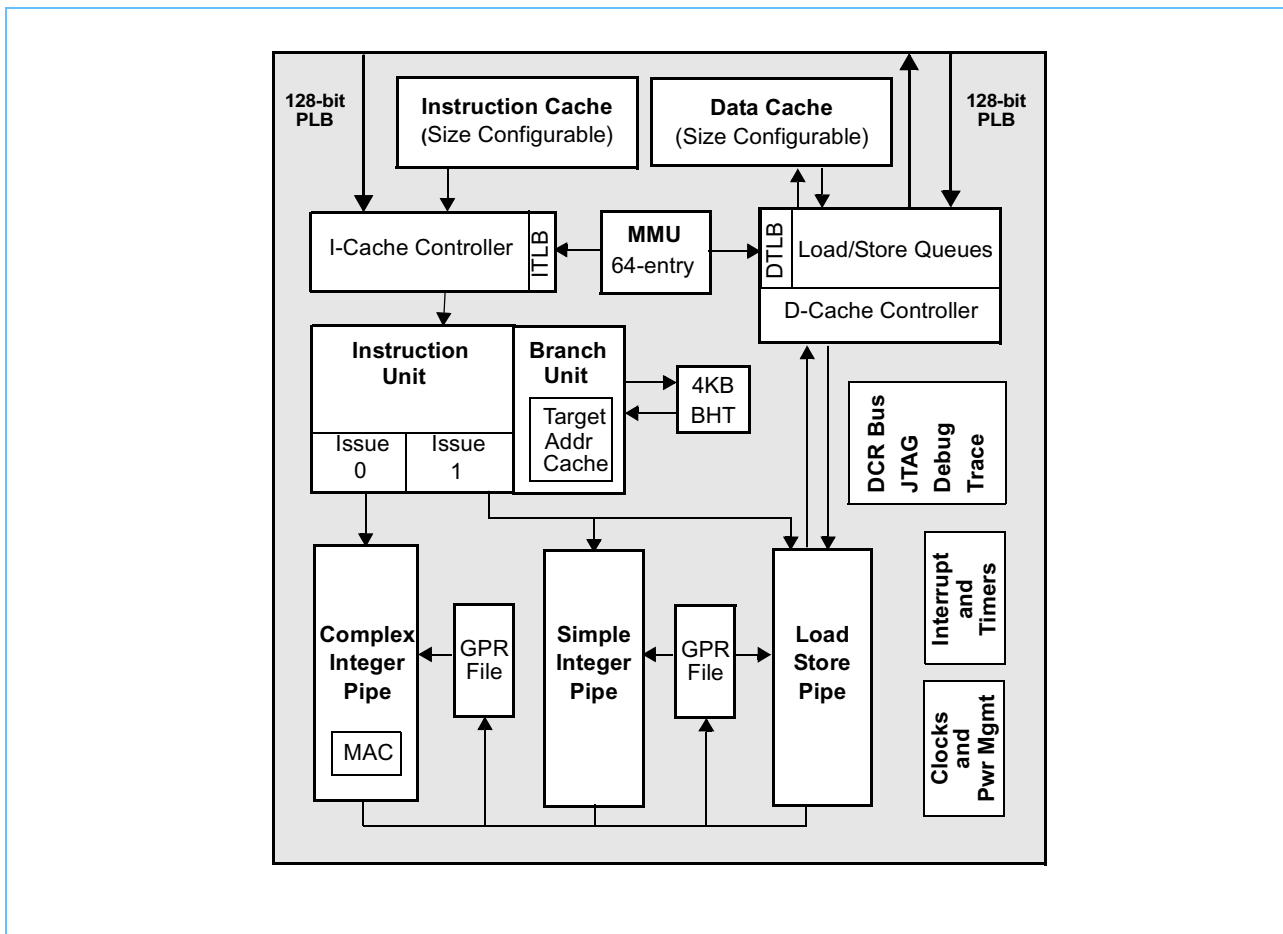
**Note:** This document differs from the Book-E architecture specification in the use of bit numbering for architected registers. Specifically, Book-E defines the full, 64-bit instruction set architecture, and thus all registers are shown as having bit numbers from 0 to 63, with bit 63 being the least significant. On the other hand, this document describes the PPC440x5 core, which is a 32-bit subset implementation of the architecture. Accordingly, all architected registers are described as being 32 bits in length, with the bits numbered from 0 to 31, and with bit 31 being the least significant. Therefore, when this document makes reference to register bit numbers from 0 to 31, they actually correspond to bits 32 to 63 of the same register in the Book-E architecture specification.

### 1.3 PPC440x5 Organization

The PPC440x5 core includes a seven-stage pipelined PowerPC core, which consists of a three stage, dual-issue instruction fetch and decode unit with attached branch unit, together with three independent, 4-stage pipelines for complex integer, simple integer, and load/store operations, respectively. The PPC440x5 core also includes a memory management unit (MMU); separate instruction and data cache units; JTAG, debug, and trace logic; and timer facilities.

Figure 1-1 illustrates the logical organization of the PPC440x5 core:

Figure 1-1. PPC440x5 Core Block Diagram



#### 1.3.1 Superscalar Instruction Unit

The instruction unit of the PPC440x5 core fetches, decodes, and issues two instructions per cycle to any combination of the three execution pipelines and/or the APU interface (see “Execution Pipelines” below, and *Auxiliary Processor Unit (APU) Port* on page 36). The instruction unit includes a branch unit which provides dynamic branch prediction using a branch history table (BHT), as well as a branch target address cache (BTAC). These mechanisms greatly improve the branch prediction accuracy and reduce the latency of taken branches, such that the target of a branch can usually be executed immediately after the branch itself, with no penalty.

### 1.3.2 Execution Pipelines

The PPC440x5 core contains three execution pipelines: complex integer, simple integer, and load/store. Each pipeline consists of four stages and can access the nine-ported (six read, three write) GPR file. In order to improve performance and avoid contention for the GPR file, there are two identical copies of it. One is dedicated to the complex integer pipeline, while the other is shared by the simple integer and the load/store pipelines.

The complex integer pipeline handles all arithmetic, logical, branch, and system management instructions (such as interrupt and TLB management, move to/from system registers, and so on). This pipeline also handles multiply and divide operations, and 24 DSP instructions that perform a variety of multiply-accumulate operations. The complex integer pipeline multiply unit can perform 32-bit  $\times$  32-bit multiply operations with single-cycle throughput and three-cycle latency; 16-bit  $\times$  32-bit multiply operations have only two-cycle latency. Divide operations take 33 cycles.

The simple integer pipeline can handle most arithmetic and logical operations which do not update the Condition Register (CR).

The load/store pipeline handles all load, store, and cache management instructions. All misaligned operations are handled in hardware, with no penalty on any operation which is contained within an aligned 16-byte region. The load/store pipeline supports all operations to both big endian and little endian data regions.

Appendix B, "PPC440x5 Core Compiler Optimizations," provides detailed information on instruction timings and performance implications in the PPC440x5 core.

### 1.3.3 Instruction and Data Cache Controllers

The PPC440x5 core provides separate instruction and data cache controllers and arrays, which allow concurrent access and minimize pipeline stalls. The storage capacity of the cache arrays, which can range from 8KB–32KB each, depends upon the implementation. Both cache controllers have 32-byte lines, and both are highly-associative, with 64-way set-associativity for 32KB and 16KB sizes, and 32-way set-associativity for the 8KB size. Both caches support parity checking on the tags and data in the memory arrays, to protect against soft errors. If a parity error is detected, the CPU will cause a machine check exception.

The PowerPC instruction set provides a rich set of cache management instructions for software-enforced coherency. The PPC440x5 implementation also provides special debug instructions that can directly read the tag and data arrays. See Chapter 4, "Instruction and Data Caches," for detailed information about the instruction and data cache controllers.

The cache controllers connect to the PLB for connection to the IBM CoreConnect system-on-a-chip environment.

#### 1.3.3.1 Instruction Cache Controller (ICC)

The ICC delivers two instructions per cycle to the instruction unit of the PPC440x5 core. The ICC also handles the execution of the PowerPC instruction cache management instructions for coherency. The ICC includes a speculative pre-fetch mechanism which can be configured to automatically pre-fetch a burst of up to three additional lines upon any fetch request which misses in the instruction cache. These speculative pre-fetches can be abandoned if the instruction execution branches away from the original instruction stream.

The ICC supports cache line locking, at either an 8-line or 16-line granularity, depending on cache size (16-line for 32KB, 8-line for 8KB and 16KB). In addition, the notion of a “transient” portion of the cache is supported, in which the cache can be configured such that only a limited portion is used for instruction cache lines from memory pages that are designated by a storage attribute from the MMU as being transient in nature. Such memory pages would contain code which is unlikely to be reused once the processor moves on to the next series of instruction lines, and thus performance may be improved by preventing each series of instruction lines from overwriting all of the “regular” code in the instruction cache.

### 1.3.3.2 Data Cache Controller (DCC)

The DCC handles all load and store data accesses, as well as the PowerPC data cache management instructions. All misaligned accesses are handled in hardware, with those accesses that are contained within a half-line (16 bytes) being handled as a single request. Load and store accesses which cross a 16-byte boundary are broken into two separate accesses by the hardware.

The DCC interfaces to the APU port to provide direct load/store access to the data cache for APU load and store operations. Such APU load and store instructions can access up to 16 bytes (one quadword) in a single cycle.

The data cache can be operated in a store-in (copy-back) or write-through manner, according to the write-through storage attribute specified for the memory page by the MMU. The DCC also supports both “store-with-allocate” and “store-without-allocate” operations, such that store operations that miss in the data cache can either “allocate” the line in the cache by reading it in and storing the new data into the cache, or alternatively bypassing the cache on a miss and simply storing the data to memory. This characteristic can also be specified on a page-by-page basis by a storage attribute in the MMU.

The DCC also supports cache line locking and “transient” data, in the same manner as the ICC (see *Instruction Cache Controller (ICC)* on page 31).

The DCC provides extensive load, store, and flush queues, such that up to three outstanding line fills and up to four outstanding load misses can be pending, and the DCC can continue servicing subsequent load and store hits in an out-of-order fashion. Store-gathering can also be performed on caching inhibited, write-through, and “without-allocate” store operations, for up to 16 contiguous bytes. Finally, each cache line has four separate “dirty” bits (one per doubleword), so that the amount of data flushed on cache line replacement can be minimized.

### 1.3.4 Memory Management Unit (MMU)

The PPC440x5 supports a flat, 36-bit (64GB) real (physical) address space. This 36-bit real address is generated by the MMU, as part of the translation process from the 32-bit effective address, which is calculated by the processor core as an instruction fetch or load/store address.

The MMU provides address translation, access protection, and storage attribute control for embedded applications. The MMU supports demand paged virtual memory and other management schemes that require precise control of logical to physical address mapping and flexible memory protection. Working with appropriate system level software, the MMU provides the following functions:

- Translation of the 32-bit effective address space into the 36-bit real address space
- Page level read, write, and execute access control
- Storage attributes for cache policy, byte order (endianness), and speculative memory access
- Software control of page replacement strategy



The translation lookaside buffer (TLB) is the primary hardware resource involved in the control of translation, protection, and storage attributes. It consists of 64 entries, each specifying the various attributes of a given page of the address space. The TLB is fully-associative; the entry for a given page can be placed anywhere in the TLB. The TLB tag and data memory arrays are parity protected against soft errors; if a parity error is detected, the CPU will cause a machine check exception.

Software manages the establishment and replacement of TLB entries. This gives system software significant flexibility in implementing a custom page replacement strategy. For example, to reduce TLB thrashing or translation delays, software can reserve several TLB entries for globally accessible static mappings. The instruction set provides several instructions for managing TLB entries. These instructions are privileged and the processor must be in supervisor state in order for them to be executed.

The first step in the address translation process is to expand the effective address into a virtual address. This is done by taking the 32-bit effective address and appending to it an 8-bit Process ID (PID), as well as a 1-bit "address space" identifier (AS). The PID value is provided by the PID register (see Chapter 5, "Memory Management"). The AS identifier is provided by the Machine State Register (MSR, see Chapter 6, "Interrupts and Exceptions," which contains separate bits for the instruction fetch address space (MSR[IS]) and the data access address space (MSR[DS]). Together, the 32-bit effective address, the 8-bit PID, and the 1-bit AS form a 41-bit virtual address. This 41-bit virtual address is then translated into the 36-bit real address using the TLB.

The MMU divides the address space (whether effective, virtual, or real) into pages. Eight page sizes (1KB, 4KB, 16KB, 64KB, 256KB, 1MB, 16MB, 256MB) are simultaneously supported, such that at any given time the TLB can contain entries for any combination of page sizes. In order for an address translation to occur, a valid entry for the page containing the virtual address must be in the TLB. An attempt to access an address for which no TLB entry exists causes an Instruction (for fetches) or Data (for load/store accesses) TLB Error exception.

To improve performance, both the instruction cache and the data cache maintain separate "shadow" TLBs. The instruction shadow TLB (ITLB) contains four entries, while the data shadow TLB (DTLB) contains eight. These shadow arrays minimize TLB contention between instruction fetch and data load/store operations. The instruction fetch and data access mechanisms only access the main 64-entry unified TLB when a miss occurs in the respective shadow TLB. The penalty for a miss in either of the shadow TLBs is three cycles. Hardware manages the replacement and invalidation of both the ITLB and DTLB; no system software action is required.

Each TLB entry provides separate user state and supervisor state read, write, and execute permission controls for the memory page associated with the entry. If software attempts to access a page for which it does not have the necessary permission, an Instruction (for fetches) or Data (for load/store accesses) Storage exception will occur.

Each TLB entry also provides a collection of storage attributes for the associated page. These attributes control cache policy (such as cachability and write-through as opposed to copy-back behavior), byte order (big endian as opposed to little endian), and enabling of speculative access for the page. In addition, a set of four, user-definable storage attributes are provided. These attributes can be used to control various system-level behaviors, such as instruction compression using IBM CodePack technology. They can also be configured to control whether data cache lines are allocated upon a store miss, and whether accesses to a given page should use the "normal" or "transient" portions of the instruction or data cache (see Chapter 4, "Instruction and Data Caches," for detailed information about these features).

Chapter 5, "Memory Management," describes the PPC440x5 MMU functions.

### 1.3.5 Timers

The PPC440x5 contains a Time Base and three timers: a Decrementer (DEC), a Fixed Interval Timer (FIT), and a Watchdog Timer. The Time Base is a 64-bit counter which gets incremented at a frequency either equal to the processor core clock rate or as controlled by a separate asynchronous timer clock input to the core. No interrupt is generated as a result of the Time Base wrapping back to zero.

The DEC is a 32-bit register that is decremented at the same rate at which the Time Base is incremented. The user loads the DEC register with a value to create the desired interval. When the register is decremented to zero, a number of actions occur: the DEC stops decrementing, a status bit is set in the Timer Status Register (TSR), and a Decrementer exception is reported to the interrupt mechanism of the PPC440x5 core. Optionally, the DEC can be programmed to reload automatically the value contained in the Decrementer Auto-Reload register (DECAR), after which the DEC resumes decrementing. The Timer Control Register (TCR) contains the interrupt enable for the Decrementer interrupt.

The FIT generates periodic interrupts based on the transition of a selected bit from the Time Base. Users can select one of four intervals for the FIT period by setting a control field in the TCR to select the appropriate bit from the Time Base. When the selected Time Base bit transitions from 0 to 1, a status bit is set in the TSR and a Fixed Interval Timer exception is reported to the interrupt mechanism of the PPC440x5 core. The FIT interrupt enable is contained in the TCR.

Similar to the FIT, the Watchdog Timer also generates a periodic interrupt based on the transition of a selected bit from the Time Base. Users can select one of four intervals for the watchdog period, again by setting a control field in the TCR to select the appropriate bit from the Time Base. Upon the first transition from 0 to 1 of the selected Time Base bit, a status bit is set in the TSR and a Watchdog Timer exception is reported to the interrupt mechanism of the PPC440x5 core. The Watchdog Timer can also be configured to initiate a hardware reset if a second transition of the selected Time Base bit occurs prior to the first Watchdog exception being serviced. This capability provides an extra measure of recoverability from potential system lock-ups.

The timer functions of the PPC440x5 core are more fully described in Chapter 7, "Timer Facilities."

### 1.3.6 Debug Facilities

The PPC440x5 debug facilities include debug modes for the various types of debugging used during hardware and software development. Also included are debug events that allow developers to control the debug process. Debug modes and debug events are controlled using debug registers in the chip. The debug registers are accessed either through software running on the processor, or through the JTAG port.

The debug modes, events, controls, and interfaces provide a powerful combination of debug facilities for hardware development tools, such as the RISCWatch™ debugger from IBM.

A brief overview of the debug modes and development tool support are provided below. Chapter 8, "Debug Facilities," provides detailed information about each debug mode and other debug resources.

#### 1.3.6.1 Debug Modes

The PPC440x5 core supports four debug modes: internal, external, real-time-trace, and debug wait. Each mode supports a different type of debug tool used in embedded systems development. Internal debug mode supports software-based ROM monitors, and external debug mode supports a hardware emulator type of debug. Real-time-trace mode uses the debug facilities to indicate events within a trace of processor execu-

tion in real time. Debug wait mode enables the processor to continue to service real-time critical interrupts while instruction execution is otherwise stopped for hardware debug. The debug modes are controlled by Debug Control Register 0 (DBCRO) and the setting of bits in the Machine State Register (MSR).

Internal debug mode supports accessing architected processor resources, setting hardware and software breakpoints, and monitoring processor status. In internal debug mode, debug events can generate debug exceptions, which can interrupt normal program flow so that monitor software can collect processor status and alter processor resources.

Internal debug mode relies on exception-handling software—running on the processor—along with an external communications path to debug software problems. This mode is used while the processor continues executing instructions and enables debugging of problems in application or operating system code. Access to debugger software executing in the processor while in internal debug mode is through a communications port on the processor board, such as a serial port or ethernet connection.

External debug mode supports stopping, starting, and single-stepping the processor, accessing architected processor resources, setting hardware and software breakpoints, and monitoring processor status. In external debug mode, debug events can architecturally “freeze” the processor. While the processor is frozen, normal instruction execution stops, and the architected processor resources can be accessed and altered using a debug tool (such as RISCWatch) attached through the JTAG port. This mode is useful for debugging hardware and low-level control software problems.

### **1.3.6.2 Development Tool Support**

The PPC440x5 provides powerful debug support for a wide range of hardware and software development tools.

The OS Open real-time operating system debugger is an example of an operating system-aware debugger, implemented using software traps.

RISCWatch is an example of a development tool that uses the external debug mode, debug events, and the JTAG port to support hardware and software development and debugging.

The RISCTrace™ feature of RISCWatch is an example of a development tool that uses the real-time trace capability of the PPC440x5.

## **1.4 Core Interfaces**

Several interfaces to the PPC440x5 core support the IBM CoreConnect on-chip system architecture, which simplifies the attachment of on-chip devices. These interfaces include:

- Processor local bus (PLB)
- Device configuration register (DCR) interface
- Auxiliary processor unit (APU) port
- JTAG, debug, and trace ports
- Interrupt interface
- Clock and power management interface

Several of these interfaces are described briefly in the sections below.

### 1.4.1 Processor Local Bus (PLB)

There are three independent 128-bit PLB interfaces to the PPC440x5 core. Each of these interfaces includes a 36-bit address bus and a 128-bit data bus. One PLB interface supports instruction cache reads, while the other two support data cache reads and writes, respectively. The frequency of each PLB interface can be independently specified, allowing an IBM CoreConnect system in which the interfaces are not all connected as part of the same PLB and in which each PLB subsystem operates at its own frequency. Each PLB interface frequency can be configured to any value such that the ratio of the processor core frequency to the PLB (core:PLB) is  $n:1$ ,  $n:2$ , or  $n:3$ , where  $n$  is any integer greater than or equal to the denominator of the ratio.

Each of the PLB interfaces supports connection to a PLB subsystem of either 32, 64, or 128 bits. The instruction and data cache controllers handle any dynamic data bus resizing which is required when the subsystem data width is less than the 128 bits of the PPC440x5 core PLB interfaces.

The data cache PLB interfaces make requests for 32-byte lines, as well as for 1 - 15 bytes within a 16-byte (quadword) aligned region. A 16-byte line request is used for quadword APU load operations to caching inhibited pages, and for quadword APU store operations to caching inhibited, write-through, or "without allocate" pages.

The instruction cache controller makes 32-byte line read requests, and also presents quadword burst read requests for up to three 32-byte lines (six quadwords), as part of its speculative line fill mechanism.

Each of the PLB interfaces fully supports the address pipelining capabilities of the PLB, and in fact can go beyond the pipeline depth and minimum latency which the PLB supports. Specifically, each interface supports up to three pipelined request/acknowledge sequences prior to performing the data transfers associated with the first request. For the data cache, if each of the requests must themselves be broken into three separate transactions (for example, for a misaligned doubleword request to a 32-bit PLB slave), then the interface actually supports up to nine outstanding request/acknowledge sequences prior to the first data transfer. Furthermore, each PLB interface tolerates a zero-cycle latency between the request and the address and data acknowledge (that is, the request, address acknowledge, and data acknowledge may all occur in the same cycle).

### 1.4.2 Device Control Register (DCR) Interface

The DCR interface provides a mechanism for the PPC440x5 core to setup other on-chip facilities. For example, programmable resources in an external bus interface unit may be configured for usage with various memory devices according to their transfer characteristics and address assignments. DCRs are accessed through the use of the PowerPC **mfocr** and **mtocr** instructions.

The interface is interlocked with control signals such that it may be connected to peripheral units that may be clocked at different frequencies from the processor core. The design allows for future expansion of the non-core facilities without changing the I/O on either the PPC440x5 core or the ASIC peripherals.

The DCR interface also allows the PPC440x5 core to communicate with peripheral devices without using the PLB interface, thereby avoiding the impact to the primary system bus bandwidth, and without additional segmentation of the useable address map.

### 1.4.3 Auxiliary Processor Unit (APU) Port

This interface provides the PPC440x5 core with the flexibility for attaching a tightly-coupled coprocessor-type macro incorporating instructions which go beyond those provided within the processor core itself. The APU port provides sufficient functionality for attachment of various coprocessor functions such as a fully-compliant

PowerPC floating point unit (single or double precision), multimedia engine, DSP, or other custom function implementing algorithms appropriate for specific system applications. The APU interface supports dual-issue pipeline designs, and can be used with macros that contain their own register files, or with simpler macros which use the CPU GPR file for source and/or target operands. APU load and store instructions can directly access the PPC440x5 data cache, with operands of up to a quadword (16 bytes) in length.

The APU interface provides the capability for a coprocessor to execute concurrently with the PPC440x5 core instructions that are not part of the PowerPC instruction set. Accordingly, areas have been reserved within the architected instruction space to allow for these customer-specific or application-specific APU instruction set extensions.

#### **1.4.4 JTAG Port**

The PPC440x5 JTAG port is enhanced to support the attachment of a debug tool such as the RISCWatch product from IBM. Through the JTAG test access port, and using the debug facilities designed into the PPC440x5 core, a debug workstation can single-step the processor and interrogate internal processor state to facilitate hardware and software debugging. The enhancements comply with the IEEE 1149.1 specification for vendor-specific extensions, and are therefore compatible with standard JTAG hardware for boundary-scan system testing.



## 2. Programming Model

The programming model of the PPC440x5 core describes how the following features and operations of the core appear to programmers:

- Storage addressing (including data types and byte ordering), starting on page 39
- Registers, starting on page 47
- Instruction classes, starting on page 52
- Instruction set, starting on page 55
- Branch processing, starting on page 62
- Integer processing, starting on page 69
- Processor control, starting on page 72
- User and supervisor state, starting on page 78
- Speculative access, starting on page 79
- Synchronization, starting on page 79

### 2.1 Storage Addressing

As a 32-bit implementation of the Book-E Enhanced PowerPC Architecture, the PPC440x5 core implements a uniform 32-bit effective address (EA) space. Effective addresses are expanded into virtual addresses and then translated to 36-bit (64GB) real addresses by the memory management unit (see *Memory Management* on page 129 for more information on the translation process). The organization of the real address space into a physical address space is system-dependent, and is described in the user's manuals for chip-level products that incorporate a PPC440x5 core.

The PPC440x5 generates an effective address whenever it executes a storage access, branch, cache management, or translation look aside buffer (TLB) management instruction, or when it fetches the next sequential instruction.

#### 2.1.1 Storage Operands

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

Data storage operands accessed by the integer load/store instructions may be bytes, halfwords, words, or—for load/store multiple and string instructions—a sequence of words or bytes, respectively. Data storage operands accessed by auxiliary processor (AP) load/store instructions can be bytes, halfwords, words, doublewords, or quadwords. The address of a storage 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, as controlled by the endian storage attribute (see *Byte Ordering* on page 42; also see *Endian (E)* on page 142 for more information on the endian storage attribute).

Operand length is implicit for each scalar storage access instruction type (that is, each storage access instruction type other than the load/store multiple and string instructions). The operand of such a scalar storage 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 storage operand is said to be *aligned* if it is aligned at its natural boundary; otherwise it is said to be *unaligned*.

Data storage operands for storage access instructions have the following characteristics.

*Table 2-1. Data Operand Definitions*

Storage Access Instruction Type	Operand Length	Addr[28:31] if aligned
Byte (or String)	8 bits	0bxxxx
Halfword	2 bytes	0bxxx0
Word (or Multiple)	4 bytes	0bxx00
Doubleword (AP only)	8 bytes	0bx000
Quadword (AP only)	16 bytes	0b0000

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

The alignment of the operand effective address of some storage access instructions may affect performance, and in some cases may cause an Alignment exception to occur. For such storage access instructions, the best performance is obtained when the storage operands are aligned. *Table 2-2* summarizes the effects of alignment on those storage access instruction types for which such effects exist. If an instruction type is not shown in the table, then there are no alignment effects for that instruction type.

*Table 2-2. Alignment Effects for Storage Access Instructions*

Storage Access Instruction Type	Alignment Effects
Integer load/store halfword	Broken into two byte accesses if crosses 16-byte boundary (EA[28:31] = 0b1111); otherwise no effect
Integer load/store word	Broken into two accesses if crosses 16-byte boundary (EA[28:31] > 0b1100); otherwise no effect
Integer load/store multiple or string	Broken into a series of 4-byte accesses until the last byte is accessed or a 16-byte boundary is reached, whichever occurs first. If bytes remain past a 16-byte boundary, resume accessing 4 bytes at a time until the last byte is accessed or the next 16-byte boundary is reached, whichever occurs first; repeat.
AP load/store halfword	Alignment exception if crosses 16-byte boundary (EA[28:31] = 0b1111); otherwise no effect (see note)
AP load/store word	Alignment exception if crosses 16-byte boundary (EA[28:31] > 0b1100); otherwise no effect (see note)
AP load/store doubleword	Alignment exception if crosses 16-byte boundary (EA[28:31] > 0b1000); otherwise no effect (see note)
AP load/store quadword	Alignment exception if crosses 16-byte boundary (EA[28:31] ≠ 0b0000); otherwise no effect

**Note:** An auxiliary processor can specify that the EA for a given AP load/store instruction must be aligned at the operand-size boundary, or alternatively, at a word boundary. If the AP so indicates this requirement and the calculated EA fails to meet it, the PPC440x5 core generates an Alignment exception. Alternatively, an auxiliary processor can specify that the EA for a given AP load/store instruction should be “forced” to be aligned, by ignoring the appropriate number of low-order EA bits and processing the AP load/store as if those bits were 0. Byte, halfword, word, doubleword, and quadword AP load/store instructions would ignore 0, 1, 2, 3, and 4 low-order EA bits, respectively.



Cache management instructions access *cache block* operands, and for the PPC440x5 core the cache block size is 32 bytes. However, the effective addresses calculated by cache management instructions are not required to be aligned on cache block boundaries. Instead, the architecture specifies that the associated low-order effective address bits (bits 27:31 for PPC440x5) are ignored during the execution of these instructions.

Similarly, the TLB management instructions access *page* operands, and—as determined by the page size—the associated low-order effective address bits are ignored during the execution of these instructions.

Instruction storage operands, on the other hand, are always four bytes long, and the effective addresses calculated by Branch instructions are therefore always word-aligned.

### 2.1.2 Effective Address Calculation

For a storage access instruction, if the sum of the effective address and the operand length exceeds the maximum effective address of  $2^{32}-1$  (that is, the storage operand itself crosses the maximum address boundary), the result of the operation is undefined, as specified by the architecture. The PPC440x5 core performs the operation as if the storage operand wrapped around from the maximum effective address to effective address 0. Software, however, should not depend upon this behavior, so that it may be ported to other implementations that do not handle this scenario in the same fashion. Accordingly, software should ensure that no data storage operands cross the maximum address boundary.

Note that since instructions are words and since the effective addresses of instructions are always implicitly on word boundaries, it is not possible for an instruction storage operand to cross any word boundary, including the maximum address boundary.

Effective address arithmetic, which calculates the starting address for storage operands, wraps around from the maximum address to address 0, for all effective address computations except next sequential instruction fetching. See *Instruction Storage Addressing Modes* on page 41 for more information on next sequential instruction fetching at the maximum address boundary.

#### 2.1.2.1 Data Storage Addressing Modes

There are two data storage addressing modes supported by the PPC440x5 core:

- Base + displacement (D-mode) 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; the low-order 32 bits of the sum form the effective address of the data storage operand.

- Base + index (X-mode) 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 0 if RA = 0; the low-order 32 bits of the sum form the effective address of the data storage operand.

#### 2.1.2.2 Instruction Storage Addressing Modes

There are four instruction storage addressing modes supported by the PPC440x5 core:

- I-form branch instructions (unconditional):

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; the low-order 32 bits of the sum form the effective address of the next instruction.

- 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; the low-order 32 bits of the sum form the effective address of the next instruction.

- Taken XL-form branch instructions:

The contents of bits 0:29 of the Link Register (LR) or the Count Register (CTR) are concatenated on the right with 0b00 to form the 32-bit effective address of the next instruction.

- Next sequential instruction fetching (including non-taken branch instructions):

The value 4 is added to the address of the current instruction to form the 32-bit effective address of the next instruction. If the address of the current instruction is 0xFFFFFFC, the PPC440x5 core wraps the next sequential instruction address back to address 0. This behavior is not required by the architecture, which specifies that the next sequential instruction address is undefined under these circumstances. Therefore, software should not depend upon this behavior, so that it may be ported to other implementations that do not handle this scenario in the same fashion. Accordingly, if software wishes to execute across this maximum address boundary and wrap back to address 0, it should place an unconditional branch at the boundary, with a displacement of 4.

In addition to the above four instruction storage addressing modes, the following behavior applies to branch instructions:

- Any branch instruction with LK=1:

The value 4 is added to the address of the current instruction and the low-order 32 bits of the result are placed into the LR. As for the similar scenario for next sequential instruction fetching, if the address of the branch instruction is 0xFFFF FFC, the result placed into the LR is architecturally undefined, although once again the PPC440x5 core wraps the LR update value back to address 0. Again, however, software should not depend on this behavior, in order that it may be ported to implementations which do not handle this scenario in the same fashion.

### 2.1.3 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 storage, 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 storage does the question of order arise.

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

For the Book-E Enhanced PowerPC Architecture, as for most current computer architectures, the smallest addressable unit of storage is the 8-bit byte. Many scalars are halfwords, words, or doublewords, which consist of groups of bytes. When a word-length scalar is moved from a register to storage, the scalar occu-

pies 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 eight 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 four 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”) eight 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 end” (most-significant end) of the scalar, considered as a binary number, comes first in storage. IBM RISC System/6000, IBM System/390, and Motorola 680x0 are examples of computer architectures using this byte ordering.

- The ordering that assigns the lowest address to the lowest-order (“right-most”) eight 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 end” (least-significant end) of the scalar, considered as a binary number, comes first in storage. The Intel x86 is an example of a processor architecture using this byte ordering.

PowerPC Book-E supports both big endian and little endian byte ordering, for both instruction and data storage accesses. Which byte ordering is used is controlled on a memory page basis by the endian (E) storage attribute, which is a field within the TLB entry for the page. The endian storage attribute is set to 0 for a big endian page, and is set to 1 for a little endian page. See *Memory Management* on page 129 for more information on memory pages, the TLB, and storage attributes, including the endian storage attribute.

### 2.1.3.1 Structure Mapping Examples

The following C language structure, *s*, contains an assortment of scalars and a character string. The comments show the value assumed to be in each structure element; these values show how the bytes comprising each structure element are mapped into storage.

```
struct {
    int a;           /* 0x1112_1314 word */
    long long b;    /* 0x2122_2324_2526_2728 doubleword */
    char *c;        /* 0x3132_3334 word */
    char d[7];      /* 'A','B','C','D','E','F','G' array of bytes */
    short e;        /* 0x5152 halfword */
    int f;          /* 0x6162_6364 word */
} s;
```

C structure mapping rules permit the use of padding (skipped bytes) to align scalars on desirable boundaries. The structure mapping examples below show each scalar aligned at its natural boundary. This alignment introduces padding of four bytes between *a* and *b*, one byte between *d* and *e*, and two bytes between *e* and *f*. The same amount of padding is present in both big endian and little endian mappings.

*Big Endian Mapping*

The big endian mapping of structure *s* follows (the data is highlighted in the structure mappings). Addresses, in hexadecimal, are below the data stored at the address. The contents of each byte, as defined in structure *s*, is shown as a (hexadecimal) number or character (for the string elements). The shaded cells correspond to padded bytes.

<b>11</b>	<b>12</b>	<b>13</b>	<b>14</b>				
0x00	0x01	0x02	0x03	0x04	0x05	0x06	0x07
<b>21</b>	<b>22</b>	<b>23</b>	<b>24</b>	<b>25</b>	<b>26</b>	<b>27</b>	<b>28</b>
0x08	0x09	0x0A	0x0B	0x0C	0x0D	0x0E	0x0F
<b>31</b>	<b>32</b>	<b>33</b>	<b>34</b>	'A'	'B'	'C'	'D'
0x10	0x11	0x12	0x13	0x14	0x15	0x16	0x17
'E'	'F'	'G'		<b>51</b>	<b>52</b>		
0x18	0x19	0x1A	0x1B	0x1C	0x1D	0x1E	0x1F
<b>61</b>	<b>62</b>	<b>63</b>	<b>64</b>				
0x20	0x21	0x22	0x23	0x24	0x25	0x26	0x27

*Little Endian Mapping*


Structure *s* is shown mapped little endian.

<b>14</b>	<b>13</b>	<b>12</b>	<b>11</b>				
0x00	0x01	0x02	0x03	0x04	0x05	0x06	0x07
<b>28</b>	<b>27</b>	<b>26</b>	<b>25</b>	<b>24</b>	<b>23</b>	<b>22</b>	<b>21</b>
0x08	0x09	0x0A	0x0B	0x0C	0x0D	0x0E	0x0F
<b>34</b>	<b>33</b>	<b>32</b>	<b>31</b>	'A'	'B'	'C'	'D'
0x10	0x11	0x12	0x13	0x14	0x15	0x16	0x17
'E'	'F'	'G'		<b>52</b>	<b>51</b>		
0x18	0x19	0x1A	0x1B	0x1C	0x1D	0x1E	0x1F
<b>64</b>	<b>63</b>	<b>62</b>	<b>61</b>				
0x20	0x21	0x22	0x23	0x24	0x25	0x26	0x27


**2.1.3.2 Instruction Byte Ordering**

PowerPC Book-E defines instructions as aligned words (four bytes) in memory. As such, instructions in a big endian program image are arranged with the most-significant byte (MSB) of the instruction word at the lowest-numbered address.

Consider the big endian mapping of instruction *p* at address 0x00, where, for example, *p* = add r7, r7, r4:



On the other hand, in a little endian mapping the same instruction is arranged with the least-significant byte (LSB) of the instruction word at the lowest-numbered address:



By the definition of PowerPC Book-E bit numbering, the most-significant byte of an instruction is the byte containing bits 0:7 of the instruction. As depicted in the instruction format diagrams (see *Instruction Formats* on page 244), this most-significant byte is the one which contains the primary opcode field (bits 0:5). Due to this difference in byte orderings, the processor must perform whatever byte reversal is required (depending on the particular byte ordering in use) in order to correctly deliver the opcode field to the instruction decoder. In the PPC440x5, this reversal is performed between the memory interface and the instruction cache, according to the value of the endian storage attribute for each memory page, such that the bytes in the instruction cache are always correctly arranged for delivery directly to the instruction decoder.

If the endian storage attribute for a memory page is reprogrammed from one byte ordering to the other, the contents of the memory page must be reloaded with program and data structures that are in the appropriate byte ordering. Furthermore, anytime the contents of instruction memory change, the instruction cache must be made coherent with the updates by invalidating the instruction cache and refetching the updated memory contents with the new byte ordering.

### 2.1.3.3 Data Byte Ordering

Unlike instruction fetches, data accesses cannot be byte-reversed between memory and the data cache. Data byte ordering in memory depends upon the data type (byte, halfword, word, and so on) of a specific data item. It is only when moving a data item of a specific type from or to an architected register (as directed by the execution of a particular storage access instruction) that it becomes known what kind of byte reversal may be required due to the byte ordering of the memory page containing the data item. Therefore, byte reversal during load or store accesses is performed between data cache (or memory, on a data cache miss, for example) and the load register target or store register source, depending on the specific type of load or store instruction (that is, byte, halfword, word, and so on).

Comparing the big endian and little endian mappings of structure *s*, as shown in *Structure Mapping Examples* on page 43, the differences between the byte locations of any data item in the structure depends upon the size of the particular data item. For example (again referring to the big endian and little endian mappings of structure *s*):

- The word *a* has its four bytes reversed within the word spanning addresses 0x00 – 0x03.
- The halfword *e* has its two bytes reversed within the halfword spanning addresses 0x1C – 0x1D.

Note that the array of bytes *d*, where each data item is a byte, is not reversed when the big endian and little endian mappings are compared. For example, the character 'A' is located at address 0x14 in both the big endian and little endian mappings.

The size of the data item being loaded or stored must be known before the processor can decide whether, and if so, how to reorder the bytes when moving them between a register and the data cache (or memory).

- For byte loads and stores, including strings, no reordering of bytes occurs, regardless of byte ordering.
- For halfword loads and stores, bytes are reversed within the halfword, for one byte order with respect to the other.

- For word loads and stores (including load/store multiple), bytes are reversed within the word, for one byte order with respect to the other.
- For doubleword loads and stores (AP loads/stores only), bytes are reversed within the doubleword, for one byte order with respect to the other.
- For quadword loads and stores (AP loads/stores only), bytes are reversed within the quadword, for one byte order with respect to the other.

Note that this mechanism applies independent of the alignment of data. In other words, when loading a multi-byte data operand with a scalar load instruction, bytes are accessed from the data cache (or memory) starting with the byte at the calculated effective address and continuing with consecutively higher-numbered bytes until the required number of bytes have been retrieved. Then, the bytes are arranged such that either the byte from the highest-numbered address (for big endian storage regions) or the lowest-numbered address (for little endian storage regions) is placed into the least-significant byte of the register. The rest of the register is filled in corresponding order with the rest of the accessed bytes. An analogous procedure is followed for scalar store instructions.

For load/store multiple instructions, each group of four bytes is transferred between memory and the register according to the procedure for a scalar load word instruction.

For load/store string instructions, the most-significant byte of the first register is transferred to or from memory at the starting (lowest-numbered) effective address, regardless of byte ordering. Subsequent register bytes (from most-significant to least-significant, and then moving into the next register, starting with the most-significant byte, and so on) are transferred to or from memory at sequentially higher-numbered addresses. This behavior for byte strings ensures that if two strings are loaded into registers and then compared, the first bytes of the strings are treated as most significant with respect to the comparison.

#### **2.1.3.4 Byte-Reverse Instructions**

PowerPC Book-E defines load/store byte-reverse instructions which can access storage which is specified as being of one byte ordering in the same manner that a regular (that is, non-byte-reverse) load/store instruction would access storage which is specified as being of the opposite byte ordering. In other words, a load/store byte-reverse instruction to a big endian memory page transfers data between the data cache (or memory) and the register in the same manner that a normal load/store would transfer the data to or from a little endian memory page. Similarly, a load/store byte-reverse instruction to a little endian memory page transfers data between the data cache (or memory) and the register in the same manner that a normal load/store would transfer the data to or from a big endian memory page.

The function of the load/store byte-reverse instructions is useful when a particular memory page contains a combination of data with both big endian and little endian byte ordering. In such an environment, the Endian storage attribute for the memory page would be set according to the predominant byte ordering for the page, and the normal load/store instructions would be used to access data operands which used this predominant byte ordering. Conversely, the load/store byte-reverse instructions would be used to access the data operands which were of the other (less prevalent) byte ordering.

Software compilers cannot typically make general use of the load/store byte-reverse instructions, so they are ordinarily used only in special, hand-coded device drivers.

## 2.2 Registers

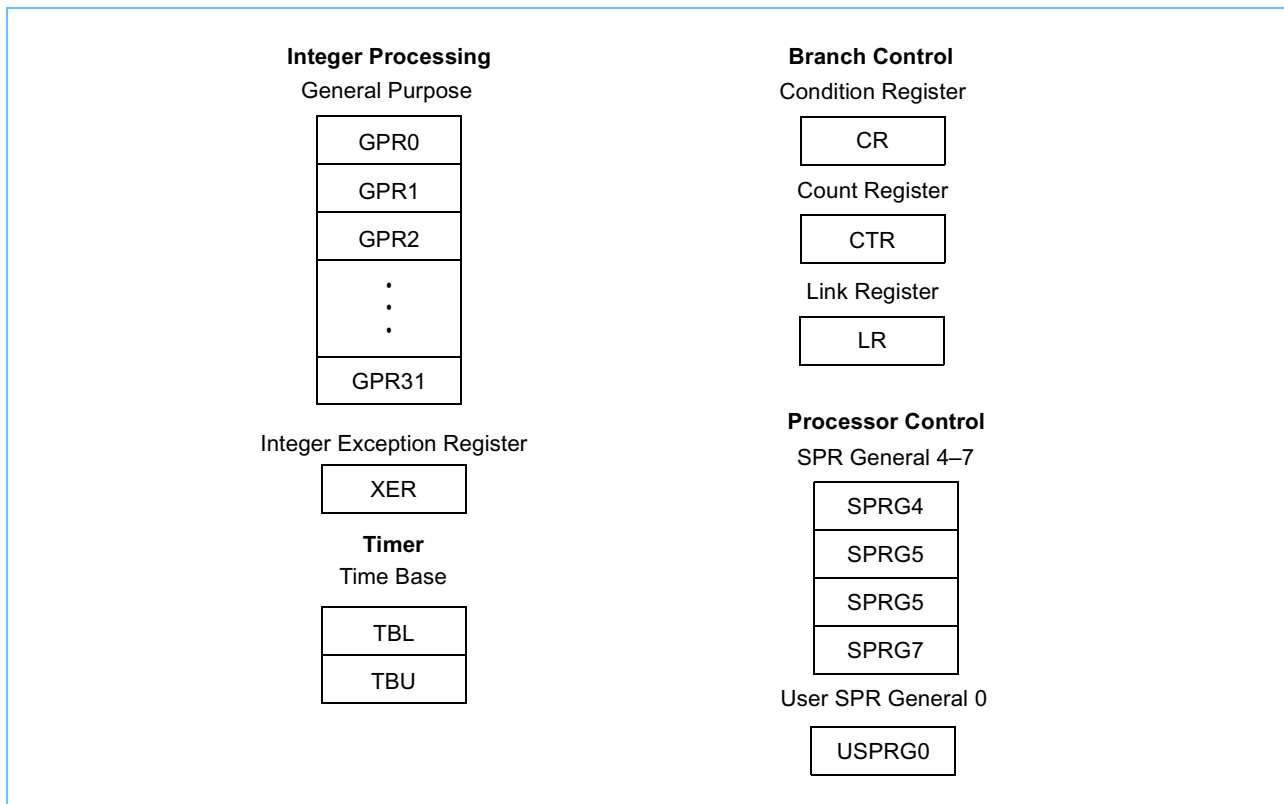
This section provides an overview of the register categories and types provided by the PPC440x5. Detailed descriptions of each of the registers are provided within the chapters covering the functions with which they are associated (for example, the cache control and cache debug registers are described in *Instruction and Data Caches* on page 93). An alphabetical summary of all registers, including bit definitions, is provided in *Register Summary* on page 441

All registers in the PPC440x5 core are architected as 32 bits wide, although certain bits in some registers are *reserved* and thus not necessarily implemented. For all registers with fields marked as reserved, these reserved fields should be written as 0 and read as *undefined*. The recommended coding practice is to perform the initial write to a register with reserved fields set to 0, and to perform all subsequent writes to the register using a read-modify-write strategy: read the register; use logical instructions to alter defined fields, leaving reserved fields unmodified; and write the register.

All of the registers are grouped into categories according to the processor functions with which they are associated. In addition, each register is classified as being of a particular *type*, as characterized by the specific instructions which are used to read and write registers of that type. Finally, most of the registers contained within the PPC440x5 core are defined by the Book-E Enhanced PowerPC Architecture, although some registers are implementation-specific and unique to the PPC440x5.

Figure 2-1 illustrates the PPC440x5 registers contained in the user programming model, that is, those registers to which access is non-privileged and which are available to both user and supervisor programs.

Figure 2-1. User Programming Model Registers



*Figure 2-2* on page 49 illustrates the PPC440x5 registers contained in the supervisor programming model, to which access is privileged and which are available to supervisor programs only. See *User and Supervisor Modes* on page 78 for more information on privileged instructions and register access, and the user and supervisor programming models.



Figure 2-2. Supervisor Programming Model Registers

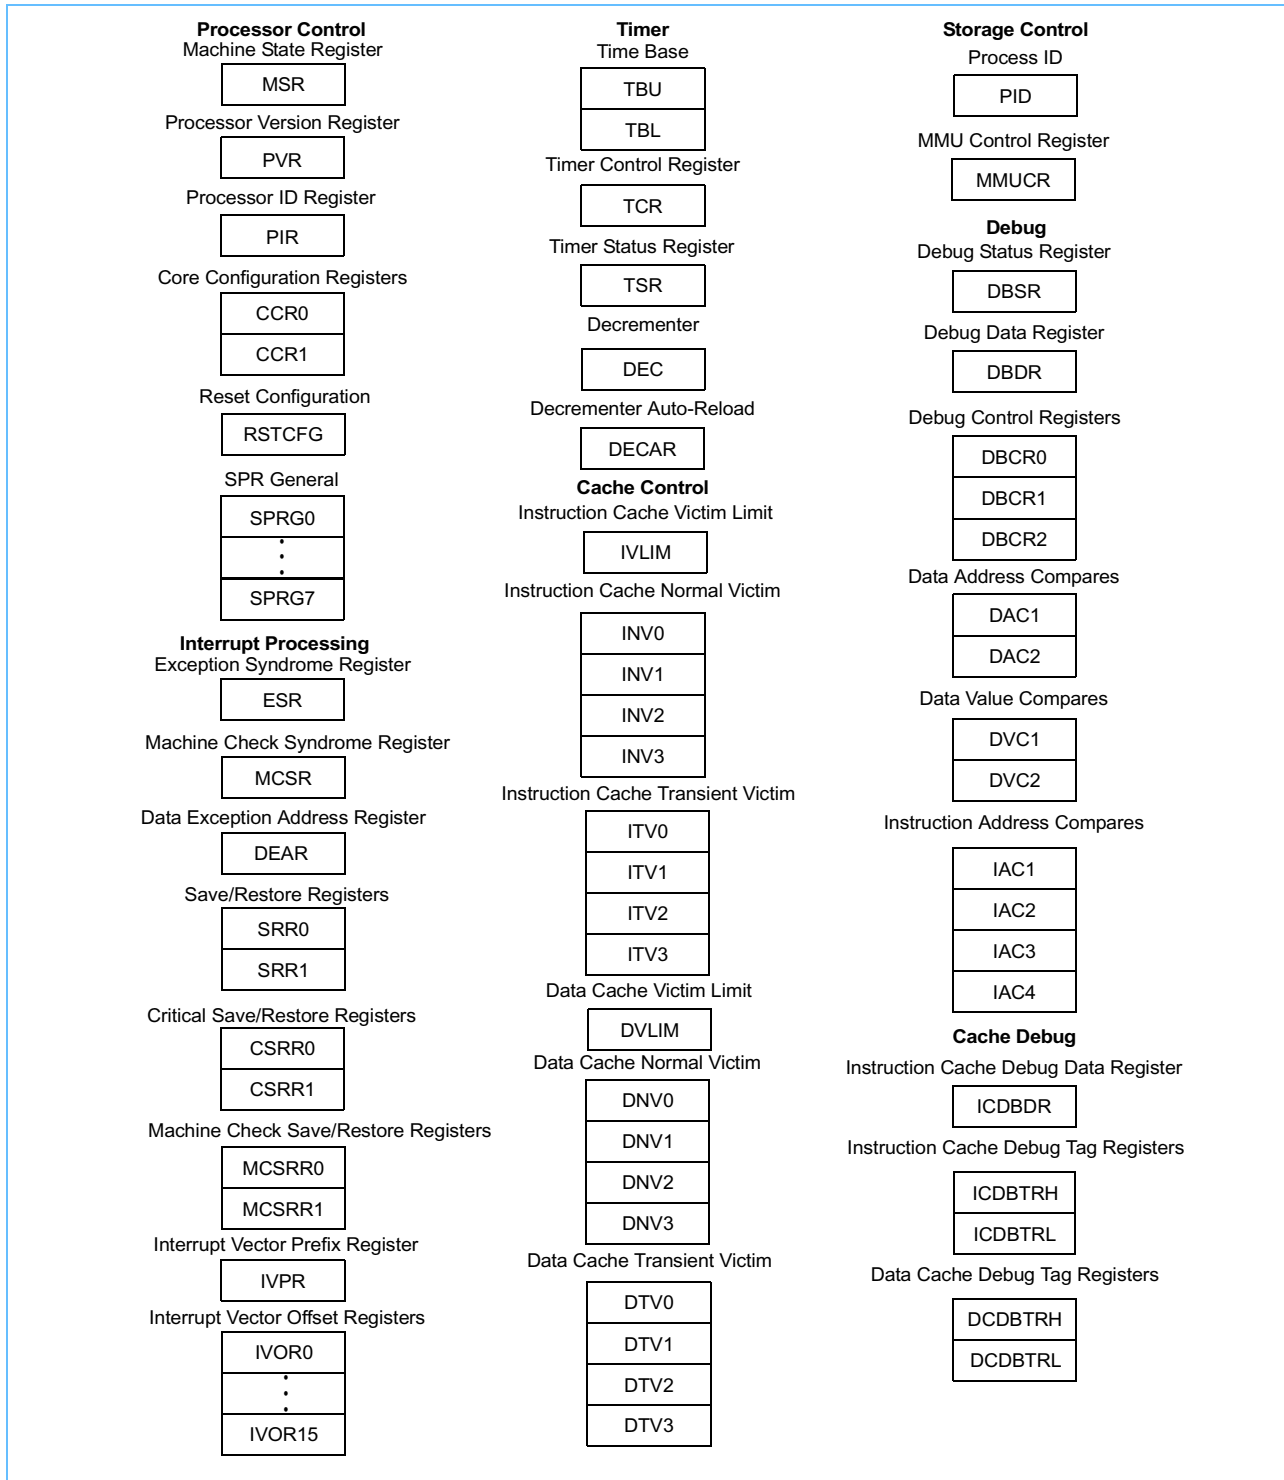


Table 2-3 lists each register category and the registers that belong to each category, along with their types and a cross-reference to the section of this document which describes them more fully. Registers that are not part of PowerPC Book-E, and are thus specific to the PPC440x5, are shown in italics in Table 2-3. Unless otherwise indicated, all registers have read/write access.

Table 2-3. Register Categories

Register Category	Register(s)	Model and Access	Type	Page
<b>Branch Control</b>	CR	User	CR	66
	CTR	User	SPR	66
	LR	User	SPR	65
<b>Cache Control</b>	<i>DNV0–DNV3</i>	Supervisor	SPR	95
	<i>DTV0–DTV3</i>	Supervisor	SPR	95
	<i>DVLIM</i>	Supervisor	SPR	97
	<i>INV0–INV3</i>	Supervisor	SPR	95
	<i>ITV0–ITV3</i>	Supervisor	SPR	95
	<i>IVLIM</i>	Supervisor	SPR	97
<b>Cache Debug</b>	<i>DCDBTRH, DCDBTRL</i>	Supervisor, read-only	SPR	123
	<i>ICDBDR, ICDBTRH, ICDBTRL</i>	Supervisor, read-only	SPR	109
<b>Debug</b>	DAC1–DAC2	Supervisor	SPR	240
	DBCR0–DBCR2	Supervisor	SPR	233
	DBDR	Supervisor	SPR	241
	DBSR	Supervisor	SPR	238
	DVC1–DVC2	Supervisor	SPR	240
	IAC1–IAC4	Supervisor	SPR	239
<b>Device Control</b>	Implemented outside core	Supervisor	DCR	52
<b>Integer Processing</b>	GPR0–GPR31	User	GPR	69
	XER	User	SPR	70
<b>Interrupt Processing</b>	CSRR0–CSRR1	Supervisor	SPR	162
	DEAR	Supervisor	SPR	164
	ESR	Supervisor	SPR	166
	IVOR0–IVOR15	Supervisor	SPR	164
	IVPR	Supervisor	SPR	165
	MCSR	Supervisor	SPR	168
	MCSRR0–MCSRR1	Supervisor	SPR	163
	SRR0–SRR1	Supervisor	SPR	161

Table 2-3. Register Categories

Register Category	Register(s)	Model and Access	Type	Page
Processor Control	CCR0	Supervisor	SPR	105
	CCR1	Supervisor	SPR	105
	MSR	Supervisor	MSR	159
	PIR, PVR	Supervisor, read-only	SPR	73
	RSTCFG	Supervisor, read-only	SPR	77
	SPRG0–SPRG3	Supervisor	SPR	73
	SPRG4–SPRG7	User, read-only; Supervisor	SPR	73
	USPRG0	User	SPR	73
Storage Control	MMUCR	Supervisor	SPR	143
	PID	Supervisor	SPR	146
Timer	DEC	Supervisor	SPR	205
	DECAR	Supervisor, write-only	SPR	205
	TBL, TBU	User read, Supervisor write	SPR	204
	TCR	Supervisor	SPR	209
	TSR	Supervisor	SPR	210

## 2.2.1 Register Types

There are five register types contained within and/or supported by the PPC440x5 core. Each register type is characterized by the instructions which are used to read and write the registers of that type. The following subsections provide an overview of each of the register types and the instructions associated with them.

### 2.2.1.1 General Purpose Registers

The PPC440x5 core contains 32 integer general purpose registers (GPRs); each contains 32 bits. Data from the data cache or memory can be loaded into GPRs using integer load instructions; the contents of GPRs can be stored to the data cache or memory using integer store instructions. Most of the integer instructions reference GPRs. The GPRs are also used as targets and sources for most of the instructions which read and write the other register types.

*Integer Processing* on page 69 provides more information on integer operations and the use of GPRs.

### 2.2.1.2 Special Purpose Registers

Special Purpose Registers (SPRs) are directly accessed using the **mtspr** and **mfspir** instructions. In addition, certain SPRs may be updated as a side-effect of the execution of various instructions. For example, the Integer Exception Register (XER) (see *Integer Exception Register (XER)* on page 70) is an SPR which is updated with arithmetic status (such as carry and overflow) upon execution of certain forms of integer arithmetic instructions.

SPRs control the use of the debug facilities, timers, interrupts, memory management, caches, and other architected processor resources. *Table 10-2* on page 443 shows the mnemonic, name, and number for each SPR, in order by SPR number. Each of the SPRs is described in more detail within the section or chapter covering the function with which it is associated. See *Table 2-3* on page 50 for a cross-reference to the associated document section for each register.

### 2.2.1.3 Condition Register

The Condition Register (CR) is a 32-bit register of its own unique type and is divided up into eight, independent 4-bit fields (CR0–CR7). The CR may be used to record certain conditional results of various arithmetic and logical operations. Subsequently, conditional branch instructions may designate a bit of the CR as one of the branch conditions (see *Branch Processing* on page 62). Instructions are also provided for performing logical bit operations and for moving fields within the CR.

See *Condition Register (CR)* on page 66 for more information on the various instructions which can update the CR.

### 2.2.1.4 Machine State Register

The Machine State Register (MSR) is a register of its own unique type that controls important chip functions, such as the enabling or disabling of various interrupt types.

The MSR can be written from a GPR using the **mtmsr** instruction. The contents of the MSR can be read into a GPR using the **mfmsr** instruction. The MSR[EE] bit can be set or cleared atomically using the **wrtee** or **wrteei** instructions. The MSR contents are also automatically saved, altered, and restored by the interrupt-handling mechanism. See *Machine State Register (MSR)* on page 159 for more detailed information on the MSR and the function of each of its bits.

### 2.2.1.5 Device Control Registers

Device Control Registers (DCRs) are on-chip registers that exist architecturally and physically outside the PPC440x5 core, and thus are not specified by the Book-E Enhanced PowerPC Architecture, nor by this user's manual for the PPC440x5 core. Rather, PowerPC Book-E simply defines the existence of the DCR address space and the instructions that access the DCRs, and does not define any particular DCRs. The DCR access instructions are **mtdcr** (move to device control register) and **mfdcr** (move from device control register), which move data between GPRs and the DCRs.

DCRs may be used to control various on-chip system functions, such as the operation of on-chip buses, peripherals, and certain processor core behaviors.

## 2.3 Instruction Classes

PowerPC Book-E architecture defines all instructions as falling into exactly one of the following four classes, as determined by the primary opcode (and the extended opcode, if any):

1. Defined
2. Allocated
3. Preserved
4. Reserved (-illegal or -nop)

### 2.3.1 Defined Instruction Class

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

One exception to this is that, for implementations (such as the PPC440x5) that only provide the 32-bit subset of PowerPC Book-E, it is not expected (and likely not even possible) that emulation of the 64-bit behavior of the defined instructions will be provided by the system.

As defined by PowerPC Book-E, any attempt to execute a defined instruction will:

- cause an Illegal Instruction exception type Program interrupt, if the instruction is not recognized by the implementation; or
- cause an Unimplemented Instruction exception type Program interrupt, if the instruction is recognized by the implementation and is not a floating-point instruction, but is not supported by the implementation; or
- cause a Floating-Point Unavailable interrupt if the instruction is recognized as a floating-point instruction, but floating-point processing is disabled; or
- cause an Unimplemented Instruction exception type Program interrupt, if the instruction is recognized as a floating-point instruction and floating-point processing is enabled, but the instruction is not supported by the implementation; or
- perform the actions described in the rest of this document, if the instruction is recognized and supported by the implementation. The architected behavior may cause other exceptions.

The PPC440x5 core recognizes and fully supports all of the instructions in the defined class, with a few exceptions. First, because the PPC440x5 is a 32-bit implementation, those operations which are defined specifically for 64-bit operation are not supported at all, and will always cause an Illegal Instruction exception type Program interrupt.

Second, instructions that are defined for floating-point processing are not supported within the PPC440x5 core, but may be implemented within an auxiliary processor and attached to the core using the AP interface. If no such auxiliary processor is attached, attempting to execute any floating-point instructions will cause an Illegal Instruction exception type Program interrupt. If an auxiliary processor which supports the floating-point instructions *is* attached, the behavior of these instructions is as defined above and as determined by the implementation details of the floating-point auxiliary processor.

Finally, there are two other defined instructions which are not supported within the PPC440x5 core. One is a TLB management instruction (**tlbiva**, TLB Invalidate Virtual Address) that is specifically intended for coherent multiprocessor systems. The other is **mfapidi** (Move From Auxiliary Processor ID Indirect), which is a special instruction intended to assist with identification of the auxiliary processors which may be attached to a particular processor implementation. Since the PPC440x5 core does not support **mfapidi**, the means of identifying the auxiliary processors in a PPC440x5 core-based system are implementation-dependent. Execution of either **tlbiva** or **mfapidi** will cause an Illegal Instruction exception type Program interrupt.

### 2.3.2 Allocated Instruction Class

This class of instructions contains a set of primary opcodes, as well as extended opcodes for certain primary opcodes. The specific opcodes are listed in *Appendix A.3* on page 543.

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

PowerPC Book-E declares that any attempt to execute an allocated instruction results in one of the following effects:

- Causes an Illegal Instruction exception type Program interrupt, if the instruction is not recognized by the implementation
- Causes an Auxiliary Processor Unavailable interrupt if the instruction is recognized by the implementation, but allocated instruction processing is disabled
- Causes an Unimplemented Instruction exception type Program interrupt, if the instruction is recognized and allocated instruction processing is enabled, but the instruction is not supported by the implementation
- Perform the actions described for the particular implementation of the allocated instruction. The implementation-dependent behavior may cause other exceptions.

In addition to supporting the defined instructions of PowerPC Book-E, the PPC440x5 also implements a number of instructions which use the allocated instruction opcodes, and thus are not part of the PowerPC Book-E architecture. *Table 2-21 on page 62* identifies the allocated instructions that are implemented within the PPC440x5 core. All of these instructions are always enabled and supported, and thus they always perform the functions defined for them within this document, and never cause Illegal Instruction, Auxiliary Processor Unavailable, nor Unimplemented Instruction exceptions.

The PPC440x5 also supports the use of *any* of the allocated opcodes by an attached auxiliary processor, except for those allocated opcodes which have been implemented within the PPC440x5 core, as mentioned above. Also, there is one other allocated opcode (primary opcode 31, secondary opcode 262) that has been implemented within the PPC440x5 core and is thus not available for use by an attached auxiliary processor. This is the opcode which was used on previous PowerPC 400 Series embedded controllers for the **icbt** (Instruction Cache Block Touch) instruction. The **icbt** instruction is now part of the defined instruction class for PowerPC Book-E, and uses a new opcode (primary opcode 31, secondary opcode 22). The PPC440x5 implements the new defined opcode, but also continues to support the previous opcode, in order to support legacy software written for earlier PowerPC 400 Series implementations. The **icbt** instruction description in *Instruction Set* on page 243 only identifies the defined opcode, although *Appendix A Instruction Summary* on page 507 includes both the defined and the allocated opcode in the table which lists all the instructions by opcode. In order to ensure portability between the PPC440x5 and future PowerPC Book-E implementations, software should take care to only use the defined opcode for **icbt**, and avoid usage of the previous opcode which is now in the allocated class.

### 2.3.3 Preserved Instruction Class

The preserved instruction class is provided to support backward compatibility with the PowerPC Architecture, and/or earlier versions of the PowerPC Book-E architecture. This instruction class includes opcodes which were defined for these previous architectures, but which are no longer defined for PowerPC Book-E.

Any attempt to execute a preserved instruction results in one of the following effects:

- Performs the actions described in the previous version of the architecture, if the instruction is recognized by the implementation
- Causes an Illegal Instruction exception type Program interrupt, if the instruction is not recognized by the implementation.

The only preserved instruction recognized and supported by the PPC440x5 is the **mftb** (Move From Time Base) opcode. This instruction was used in the the PowerPC Architecture to read the Time Base Upper (TBU) and Time Base Lower (TBL) registers. PowerPC Book-E architecture instead defines TBU and TBL as Special Purpose Registers (SPRs), and thus the **mfspr** (Move From Special Purpose Register) instruction is used to read them. In order to enable legacy time base management software to be run on the PPC440x5, the core also supports the preserved opcode of **mftb**. However, the **mftb** instruction is not included in the various sections of this document that describe the implemented instructions, and software should take care to use the currently architected mechanism of **mfspr** to read the time base registers, in order to guarantee portability between the PPC440x5 and future implementations of PowerPC Book-E.

On the other hand, *Appendix A Instruction Summary* on page 507 does identify the **mftb** instruction as an implemented preserved opcode in the table which lists all the instructions by opcode.

### 2.3.4 Reserved Instruction Class

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

Reserved instructions are available for future versions of PowerPC Book-E architecture. That is, future versions of PowerPC 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.

Any attempt to execute a reserved-illegal instruction will cause an Illegal Instruction exception type Program interrupt on implementations (such as the PPC440x5) that conform to the current version of PowerPC Book-E. Reserved-illegal instructions are, therefore, available for future extensions to PowerPC Book-E that would affect architected state. Such extensions might include new forms of integer or floating-point arithmetic instructions, or new forms of load or store instructions that affect architected registers or the contents of memory.

Any attempt to execute a reserved-nop instruction, on the other hand, either has no effect (that is, is treated as a no-operation instruction), or causes an Illegal Instruction exception type Program interrupt, on implementations (such as the PPC440x5) that conform to the current version of PowerPC Book-E. Because implementations are typically expected to treat reserved-nop instructions as true no-ops, these instruction opcodes are thus available for future extensions to PowerPC Book-E which have no effect on architected state. Such extensions might include performance-enhancing hints, such as new forms of cache touch instructions. Software would be able to take advantage of the functionality offered by the new instructions, and still remain backwards-compatible with implementations of previous versions of PowerPC Book-E.

The PPC440x5 implements all of the reserved-nop instruction opcodes as true no-ops. The specific reserved-nop opcodes are listed in *Appendix A.5* on page 544

## 2.4 Implemented Instruction Set Summary

This section provides an overview of the various types and categories of instructions implemented within the PPC440x5. In addition, *Instruction Set* on page 243 provides a complete alphabetical listing of every implemented instruction, including its register transfer language (RTL) and a detailed description of its operation. Also, *Appendix A Instruction Summary* on page 507 lists each implemented instruction alphabetically (and by opcode) along with a short-form description and its extended mnemonic(s).

Table 2-4 summarizes the PPC440x5 instruction set by category. Instructions within each category are described in subsequent sections.

Table 2-4. Instruction Categories

Category	Subcategory	Instruction Types
<b>Integer</b>	Integer Storage Access	load, store
	Integer Arithmetic	add, subtract, multiply, divide, negate
	Integer Logical	and, andc, or, orc, xor, nand, nor, xnor, extend sign, count leading zeros
	Integer Compare	compare, compare logical
	Integer Select	select operand
	Integer Trap	trap
	Integer Rotate	rotate and insert, rotate and mask
	Integer Shift	shift left, shift right, shift right algebraic
<b>Branch</b>		branch, branch conditional, branch to link, branch to count
<b>Processor Control</b>	Condition Register Logical	crand, crandc, cror, crorc, crnand, crnor, crxor, crxnor
	Register Management	move to/from SPR, move to/from DCR, move to/from MSR, write to external interrupt enable bit, move to/from CR
	System Linkage	system call, return from interrupt, return from critical interrupt, return from machine check interrupt
	Processor Synchronization	instruction synchronize
<b>Storage Control</b>	Cache Management	data allocate, data invalidate, data touch, data zero, data flush, data store, instruction invalidate, instruction touch
	TLB Management	read, write, search, synchronize
	Storage Synchronization	memory synchronize, memory barrier
<b>Allocated</b>	Allocated Arithmetic	multiply-accumulate, negative multiply-accumulate, multiply halfword
	Allocated Logical	detect left-most zero byte
	Allocated Cache Management	data congruence-class invalidate, instruction congruence-class invalidate
	Allocated Cache Debug	data read, instruction read

## 2.4.1 Integer Instructions

Integer instructions transfer data between memory and the GPRs, and perform various operations on the GPRs. This category of instructions is further divided into seven sub-categories, described below.

### 2.4.1.1 Integer Storage Access Instructions

Integer storage access instructions load and store data between memory and the GPRs. These instructions operate on bytes, halfwords, and words. Integer storage access instructions also support loading and storing multiple registers, character strings, and byte-reversed data, and loading data with sign-extension.

Table 2-5 shows the integer storage access instructions in the PPC440x5. In the table, the syntax “[u]” indicates that the instruction has both an “update” form (in which the RA addressing register is updated with the calculated address) and a “non-update” form. Similarly, the syntax “[x]” indicates that the instruction has both



an “indexed” form (in which the address is formed by adding the contents of the RA and RB GPRs) and a “base + displacement” form (in which the address is formed by adding a 16-bit signed immediate value (specified as part of the instruction) to the contents of GPR RA. See the detailed instruction descriptions in *Instruction Set* on page 243.

Table 2-5. Integer Storage Access Instructions

Loads				Stores			
Byte	Halfword	Word	Multiple/String	Byte	Halfword	Word	Multiple/String
lbz[u][x]	lha[u][x] lhbrx lhz[u][x]	lwarx lwbrx lwz[u][x]	lmw lswi lswx	stb[u][x]	sth[u][x] sthbrx	stw[u][x] stwbrx stwcx.	stmw stswi stswx

### 2.4.1.2 Integer Arithmetic Instructions

Arithmetic operations are performed on integer or ordinal operands stored in registers. Instructions that perform operations on two operands are defined in a three-operand format; an operation is performed on the operands, which are stored in two registers. The result is placed in a third register. Instructions that perform operations on one operand are defined in a two-operand format; the operation is performed on the operand in a register and the result is placed in another register. Several instructions also have immediate formats in which one of the source operands is a field in the instruction.

Most integer arithmetic instructions have versions that can update CR[CR0] and/or XER[SO, OV] (Summary Overflow, Overflow), based on the result of the instruction. Some integer arithmetic instructions also update XER[CA] (Carry) implicitly. See *Integer Processing* on page 69 for more information on how these instructions update the CR and/or the XER.

Table 2-6 lists the integer arithmetic instructions in the PPC440x5. In the table, the syntax “[o]” indicates that the instruction has both an “o” form (which updates the XER[SO,OV] fields) and a “non-o” form. Similarly, the syntax “[.]” indicates that the instruction has both a “record” form (which updates CR[CR0]) and a “non-record” form.

Table 2-6. Integer Arithmetic Instructions

Add	Subtract	Multiply	Divide	Negate
add[o][.] addc[o][.] adde[o][.] addi addic[.] addis addme[o][.] addze[o][.]	subf[o][.] subfc[o][.] subfe[o][.] subfic subfme[o][.] subfze[o][.]	mulhw[.] mulhwu[.] mulli mullw[o][.]	divw[o][.] divwu[o][.]	neg[o][.]

### 2.4.1.3 Integer Logical Instructions

Table 2-7 lists the integer logical instructions in the PPC440x5. See *Integer Arithmetic Instructions* on page 57 for an explanation of the “[.]” syntax.

Table 2-7. Integer Logical Instructions

And	And with complement	Nand	Or	Or with complement	Nor	Xor	Equivalence	Extend sign	Count leading zeros
<b>and[.]</b> <b>andi.</b> <b>andis.</b>	<b>andc[.]</b>	<b>nand[.]</b>	<b>or[.]</b> <b>ori</b> <b>oris</b>	<b>orc[.]</b>	<b>nor[.]</b>	<b>xor[.]</b> <b>xori</b> <b>xoris</b>	<b>eqv[.]</b>	<b>extsb[.]</b> <b>extsh[.]</b>	<b>cntlzw[.]</b>

### 2.4.1.4 Integer Compare Instructions

These instructions perform arithmetic or logical comparisons between two operands and update the CR with the result of the comparison.

Table 2-8 lists the integer compare instructions in the PPC440x5.

Table 2-8. Integer Compare Instructions

Arithmetic	Logical
<b>cmp</b> <b>cmpi</b>	<b>cmpl</b> <b>cmpli</b>

### 2.4.1.5 Integer Trap Instructions

Table 2-9 lists the integer trap instructions in the PPC440x5.

Table 2-9. Integer Trap Instructions

Trap
<b>tw</b> <b>twi</b>

### 2.4.1.6 Integer Rotate Instructions

These instructions rotate operands stored in the GPRs. Rotate instructions can also mask rotated operands.

Table 2-10 lists the rotate instructions in the PPC440x5. See *Integer Arithmetic Instructions* on page 57 for an explanation of the “[.]” syntax.

Table 2-10. Integer Rotate Instructions

Rotate and Insert	Rotate and Mask
<b>rlwimi[.]</b>	<b>rlwinm[.]</b> <b>rlwnm[.]</b>

### 2.4.1.7 Integer Shift Instructions

Table 2-11 lists the integer shift instructions in the PPC440x5. Note that the shift right algebraic instructions implicitly update the XER[CA] field. See *Integer Arithmetic Instructions* on page 57 for an explanation of the “[.]” syntax.

Table 2-11. Integer Shift Instructions

Shift Left	Shift Right	Shift Right Algebraic
<b>slw</b> [.]	<b>srw</b> [.]	<b>sraw</b> [.] <b>srawi</b> [.]

### 2.4.1.8 Integer Select Instruction

Table 2-12 lists the integer select instruction in the PPC440x5. The RA operand is 0 if the RA field of the instruction is 0, or is the contents of GPR[RA] otherwise.

Table 2-12. Integer Select Instruction

Integer Select
<b>isel</b>

## 2.4.2 Branch Instructions

These instructions unconditionally or conditionally branch to an address. Conditional branch instructions can test condition codes set in the CR by a previous instruction and branch accordingly. Conditional branch instructions can also decrement and test the Count Register (CTR) as part of branch determination, and can save the return address in the Link Register (LR). The target address for a branch can be a displacement from the current instruction address or an absolute address, or contained in the LR or CTR.

See *Branch Processing* on page 62 for more information on branch operations.

Table 2-13 lists the branch instructions in the PPC440x5. In the table, the syntax “[l]” indicates that the instruction has both a “link update” form (which updates LR with the address of the instruction after the branch) and a “non-link update” form. Similarly, the syntax “[a]” indicates that the instruction has both an “absolute address” form (in which the target address is formed directly using the immediate field specified as part of the instruction) and a “relative” form (in which the target address is formed by adding the specified immediate field to the address of the branch instruction).

Table 2-13. Branch Instructions

Branch
<b>b</b> [l][a]
<b>bc</b> [l][a]
<b>bcctr</b> [l]
<b>bclr</b> [l]

### 2.4.3 Processor Control Instructions

Processor control instructions manipulate system registers, perform system software linkage, and synchronize processor operations. The instructions in these three sub-categories of processor control instructions are described below.

### 2.4.3.1 Condition Register Logical Instructions

These instructions perform logical operations on a specified pair of bits in the CR, placing the result in another specified bit. The benefit of these instructions is that they can logically combine the results of several comparison operations without incurring the overhead of conditional branching between each one. Software performance can significantly improve if multiple conditions are tested at once as part of a branch decision.

Table 2-14 lists the condition register logical instructions in the PPC440x5.

Table 2-14. Condition Register Logical Instructions

<b>crand</b>	<b>crnor</b>
<b>crandc</b>	<b>cror</b>
<b>creqv</b>	<b>crorc</b>
<b>crnand</b>	<b>crxor</b>

### 2.4.3.2 Register Management Instructions

These instructions move data between the GPRs and control registers in the PPC440x5.

Table 2-15 lists the register management instructions in the PPC440x5.

Table 2-15. Register Management Instructions

CR	DCR	MSR	SPR
<b>mcrf</b>		<b>mfmsr</b>	
<b>mcrxr</b>	<b>mfdcr</b>	<b>mtmsr</b>	<b>mfspr</b>
<b>mfcr</b>	<b>mtdcr</b>	<b>wrtee</b>	<b>mtspr</b>
<b>mtrf</b>		<b>wrteei</b>	

### 2.4.3.3 System Linkage Instructions

These instructions invoke supervisor software level for system services, and return from interrupts.

Table 2-16 lists the system linkage instructions in the PPC440x5.

Table 2-16. System Linkage Instructions

<b>rfi</b>
<b>rfci</b>
<b>rfmci</b>
<b>sc</b>

### 2.4.3.4 Processor Synchronization Instruction

The processor synchronization instruction, **isync**, forces the processor to complete all instructions preceding the **isync** before allowing any context changes as a result of any instructions that follow the **isync**. Additionally, all instructions that follow the **isync** will execute within the context established by the completion of all the instructions that precede the **isync**. See *Synchronization* on page 79 for more information on the synchronizing effect of **isync**.

Table 2-17 shows the processor synchronization instruction in the PPC440x5.

Table 2-17. Processor Synchronization Instruction

<b>isync</b>
--------------

## 2.4.4 Storage Control Instructions

These instructions manage the instruction and data caches and the TLB of the PPC440x5 core. Instructions are also provided to synchronize and order storage accesses. The instructions in these three sub-categories of storage control instructions are described below.

### 2.4.4.1 Cache Management Instructions

These instructions control the operation of the data and instruction caches. Instructions are provided to fill, flush, invalidate, or zero data cache blocks, where a block is defined as a 32-byte cache line. Instructions are also provided to fill or invalidate instruction cache blocks.

Table 2-18 lists the cache management instructions in the PPC440x5.

Table 2-18. Cache Management Instructions

Data Cache	Instruction Cache
dcba dcbf dcbi dcbst dcbt dcbtst dcbz	icbi icbt

### 2.4.4.2 TLB Management Instructions

The TLB management instructions read and write entries of the TLB array, and search the TLB array for an entry which will translate a given virtual address. There is also an instruction for synchronizing TLB updates with other processors, but since the PPC440x5 core is intended for use in uni-processor environments, this instruction performs no operation on the PPC440x5.

Table 2-19 lists the TLB management instructions in the PPC440x5. See *Integer Arithmetic Instructions* on page 57 for an explanation of the “[.]” syntax.

Table 2-19. TLB Management Instructions

tlbre tlbsx[.] tlbsync tlbwe
---------------------------------------

### 2.4.4.3 Storage Synchronization Instructions

The storage synchronization instructions allow software to enforce ordering amongst the storage accesses caused by load and store instructions, which by default are “weakly-ordered” by the processor. “Weakly-ordered” means that the processor is architecturally permitted to perform loads and stores generally out-of-order with respect to their sequence within the instruction stream, with some exceptions. However, if a storage synchronization instruction is executed, then all storage accesses prompted by instructions preceding the synchronizing instruction must be performed before any storage accesses prompted by instructions which come after the synchronizing instruction. See *Synchronization* on page 79 for more information on storage synchronization.

Table 2-17 shows the storage synchronization instructions in the PPC440x5.

Table 2-20. Storage Synchronization Instructions

msync
mbar

### 2.4.5 Allocated Instructions

These instructions are not part of the PowerPC Book-E architecture, but they are included as part of the PPC440x5 core. Architecturally, they are considered allocated instructions, as they use opcodes which are within the allocated class of instructions, which the PowerPC Book-E architecture identifies as being available for implementation-dependent and/or application-specific purposes. However, all of the allocated instructions which are implemented within the PPC440x5 core are “standard” for IBM’s family of PowerPC embedded controllers, and are not unique to the PPC440x5.

The allocated instructions implemented within the PPC440x5 are divided into four sub-categories, and are shown in Table 2-21. See *Integer Arithmetic Instructions* on page 57 for an explanation of the “[.]” and “[o]” syntax.

Table 2-21. Allocated Instructions

Arithmetic			Logical	Cache Management	Cache Debug
Multiply-Accumulate	Negative Multiply-Accumulate	Multiply Halfword			
macchw[o][.] macchws[o][.] macchwsu[o][.] macchwu[o][.] machhw[o][.] machhws[o][.] machhwsu[o][.] machhwu[o][.] machw[o][.] machws[o][.] machwsu[o][.] machwu[o][.]	nmacchw[o][.] nmacchws[o][.] nmachhw[o][.] nmachhws[o][.] nmaclhw[o][.] nmaclhws[o][.]	mulchw[.] mulchwu[.] mulhhw[.] mulhhwu[.] mullhw[.] mullhwu[.]	dImzb[.]	dccci iccci	dcread icread

## 2.5 Branch Processing

The four branch instructions provided by PPC440x5 are summarized in Table 2.4.2 on page 59. In addition, each of these instructions is described in detail in *Instruction Set* on page 243. The following sections provide additional information on branch addressing, instruction fields, prediction, and registers.

### 2.5.1 Branch Addressing

The branch instruction (**b[I][a]**) specifies the displacement of the branch target address as a 26-bit value (the 24-bit LI field right-extended with 0b00). This displacement is regarded as a signed 26-bit number covering an address range of ±32MB. Similarly, the branch conditional instruction (**bc[I][a]**) specifies the displacement as a 16-bit value (the 14-bit BD field right-extended with 0b00). This displacement covers an address range of ±32KB.

For the relative form of the branch and branch conditional instructions (**b[l]** and **bc[l]**), with instruction field AA = 0), the target address is the address of the branch instruction itself (the Current Instruction Address, or CIA) plus the signed displacement. This address calculation is defined to “wrap around” from the maximum effective address (0xFFFFFFFF) to 0x0000 0000, and vice-versa.

For the absolute form of the branch and branch conditional instructions (**ba[l]** and **bca[l]**), with instruction field AA = 1), the target address is the sign-extended displacement. This means that with absolute forms of the branch and branch conditional instructions, the branch target can be within the first or last 32MB or 32KB of the address space, respectively.

The other two branch instructions, **bclr** (branch conditional to LR) and **bcctr** (branch conditional to CTR), do not use absolute nor relative addressing. Instead, they use *indirect* addressing, in which the target of the branch is specified indirectly as the contents of the LR or CTR.

### 2.5.2 Branch Instruction BI Field

Conditional branch instructions can optionally test one bit of the CR, as indicated by instruction field BO[0] (see BO field description below). The value of instruction field BI specifies the CR bit to be tested (0-31). The BI field is ignored if BO[0] = 1. The branch (**b[l][a]**) instruction is by definition unconditional, and hence does not have a BI instruction field. Instead, the position of this field is part of the LI displacement field.

### 2.5.3 Branch Instruction BO Field

The BO field specifies the condition under which a conditional branch is taken, and whether the branch decrements the CTR. The branch (**b[l][a]**) instruction is by definition unconditional, and hence does not have a BO instruction field. Instead, the position of this field is part of the LI displacement field.

Conditional branch instructions can optionally test one bit in the CR. This option is selected when BO[0] = 0; if BO[0] = 1, the CR does not participate in the branch condition test. If the CR condition option is selected, the condition is satisfied (branch can occur) if the CR bit selected by the BI instruction field matches BO[1].

Conditional branch instructions can also optionally decrement the CTR by one, and test whether the decremented value is 0. This option is selected when BO[2] = 0; if BO[2] = 1, the CTR is not decremented and does not participate in the branch condition test. If CTR decrement option is selected, BO[3] specifies the condition that must be satisfied to allow the branch to be taken. If BO[3] = 0, CTR ≠ 0 is required for the branch to occur. If BO[3] = 1, CTR = 0 is required for the branch to occur.

Table 2-22 summarizes the usage of the bits of the BO field. BO[4] is further discussed in *Branch Prediction* on page 64

Table 2-22. BO Field Definition

BO Bit	Description
BO[0]	CR Test Control 0 Test CR bit specified by BI field for value specified by BO[1] 1 Do not test CR
BO[1]	CR Test Value 0 If BO[0] = 0, test for CR[BI] = 0. 1 If BO[0] = 0, test for CR[BI] = 1.
BO[2]	CTR Decrement and Test Control 0 Decrement CTR by one and test whether the decremented CTR satisfies the condition specified by BO[3]. 1 Do not decrement CTR, do not test CTR.

Table 2-22. BO Field Definition (continued)

BO Bit	Description
BO[3]	CTR Test Value 0 If BO[2] = 0, test for decremented CTR ≠ 0. 1 If BO[2] = 0, test for decremented CTR = 0.
BO[4]	Branch Prediction Reversal 0 Apply standard branch prediction. 1 Reverse the standard branch prediction.

Table 2-23 lists specific BO field contents, and the resulting actions; z represents a mandatory value of zero, and y is a branch prediction option discussed in *Branch Prediction* on page 64

Table 2-23. BO Field Examples

BO Value	Description
0000y	Decrement the CTR, then branch if the decremented CTR ≠ 0 and CR[BI]=0.
0001y	Decrement the CTR, then branch if the decremented CTR = 0 and CR[BI] = 0.
001zy	Branch if CR[BI] = 0.
0100y	Decrement the CTR, then branch if the decremented CTR ≠ 0 and CR[BI] = 1.
0101y	Decrement the CTR, then branch if the decremented CTR=0 and CR[BI] = 1.
011zy	Branch if CR[BI] = 1.
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.

### 2.5.4 Branch Prediction

Conditional branches might be taken or not taken; if taken, instruction fetching is re-directed to the target address. If the branch is not taken, instruction fetching simply falls through to the next sequential instruction. The PPC440x5 core attempts to predict whether or not a branch is taken before all information necessary to determine the branch direction is available. This action is called *branch prediction*. The core can then prefetch instructions down the predicted path. If the prediction is correct, performance is improved because the branch target instruction is available immediately, instead of having to wait until the branch conditions are resolved. If the prediction is incorrect, then the prefetched instructions (which were fetched from addresses down the “wrong” path of the branch) must be discarded, and new instructions fetched from the correct path.

The PPC440x5 core combines the static prediction mechanism defined by PowerPC Book-E, together with a dynamic branch prediction mechanism, in order to provide correct branch prediction as often as possible. The dynamic branch prediction mechanism is an implementation optimization, and is not part of the architecture, nor is it visible to the programming model. *Appendix B PPC440x5 Core Compiler Optimizations* on page 553 provides additional information on the dynamic branch prediction mechanism.

The static branch prediction mechanism enables software to designate the “preferred” branch prediction via bits in the instruction encoding. The “default” static branch prediction for conditional branches is as follows:

$$\text{Predict that the branch is to be taken if } ((\text{BO}[0] \wedge \text{BO}[2]) \vee s) = 1$$

where *s* is bit 16 of the instruction (the sign bit of the displacement for all **bc** forms, and zero for all **bclr** and **bcctr** forms). In other words, conditional branches are predicted taken if they are “unconditional” (i.e., they do not test the CR nor the CTR decrement, and are always taken), or if their branch displacement is “negative” (i.e., the branch is branching “backwards” from the current instruction address). The standard prediction for



this case derives from considering the relative form of **bc**, often used at the end of loops to control the number of times that a loop is executed. The branch is taken each time the loop is executed except the last, so it is best if the branch is predicted taken. The branch target is the beginning of the loop, so the branch displacement is negative and  $s = 1$ . Because this situation is most common, a branch is taken if  $s = 1$ .

If branch displacements are positive,  $s = 0$ , then the branch is predicted not taken. Also, if the branch instruction is any form of **bclr** or **bcctr** except the “unconditional” form, then  $s = 0$ , and the branch is predicted not taken.

There is a peculiar consequence of this prediction algorithm for the absolute forms of **bc** (**bca** and **bcla**). As described in *Branch Addressing* on page 62, if  $s = 1$ , the branch target is in high memory. If  $s = 0$ , the branch target is in low memory. Because these are absolute-addressing forms, there is no reason to treat high and low memory differently. Nevertheless, for the high memory case the standard prediction is taken, and for the low memory case the standard prediction is not taken.

Another bit in the BO field allows software further control over branch prediction. Specifically, BO[4] is the *prediction reversal bit*. If BO[4] = 0, the default prediction is applied. If BO[4] = 1, the reverse of the default prediction is applied. For the cases in *Table 2-23* where BO[4] =  $y$ , software can reverse the default prediction by setting  $y$  to 1. This should only be done when the default prediction is likely to be wrong. Note that for the “branch always” condition, reversal of the default prediction is not allowed, as BO[4] is designated as  $z$  for this case, meaning the bit must be set to 0 or the instruction form is invalid.

## 2.5.5 Branch Control Registers

There are three registers in the PPC440x5 which are associated with branch processing, and they are described in the following sections.

### 2.5.5.1 Link Register (LR)

The LR is written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**. The LR can also be updated by the “link update” form of branch instructions (instruction field LK = 1). Such branch instructions load the LR with the address of the instruction following the branch instruction (4 + address of the branch instruction). Thus, the LR contents can be used as a return address for a subroutine that was entered using a link update form of branch. The **bclr** instruction uses the LR in this fashion, enabling indirect branching to any address.

When being used as a return address by a **bclr** instruction, bits 30:31 of the LR are ignored, since all instruction addresses are on word boundaries.

Access to the LR is non-privileged.




Figure 2-3. Link Register (LR)



### 2.5.5.2 Count Register (CTR)

The CTR is written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**. The CTR contents can be used as a loop count that gets decremented and tested by conditional branch instructions that specify count decrement as one of their branch conditions (instruction field BO[2] = 0). Alternatively, the CTR contents can specify a target address for the **bcctr** instruction, enabling indirect branching to any address.

Access to the CTR is non-privileged.




Figure 2-4. Count Register (CTR)

0:31	Count	Used as count for branch conditional with decrement instructions, or as target address for <b>bcctr</b> instructions
------	-------	--

### 2.5.5.3 Condition Register (CR)

The CR is used to record certain information (“conditions”) related to the results of the various instructions which are enabled to update the CR. A bit in the CR may also be selected to be tested as part of the condition of a conditional branch instruction.

The CR is organized into eight 4-bit fields (CR0–CR7), as shown in *Figure 2-5*. *Table 2-24* lists the instructions which update the CR.

Access to the CR is non-privileged.




Figure 2-5. Condition Register (CR)

0:3	CR0	Condition Register Field 0
4:7	CR1	Condition Register Field 1
8:11	CR2	Condition Register Field 2
12:15	CR3	Condition Register Field 3
16:19	CR4	Condition Register Field 4
20:23	CR5	Condition Register Field 5
24:27	CR6	Condition Register Field 6
28:31	CR7	Condition Register Field 7

Table 2-24. CR Updating Instructions

Integer						Processor Control	Storage Control	Auxiliary Processor
Storage Access	Arithmetic	Logical	Compare	Rotate	Shift	CR-Logical and Register Management	TLB Mgmt.	Arithmetic and Logical
stwcx.	add.[o] addc.[o] adde.[o] addic. addme.[o] addze.[o]  subf.[o] subfc.[o] subfe.[o] subfme.[o] subfze.[o]  mulhw. mulhwu. mullw.[o]  divw.[o] divwu.[o]  neg.[o]	and. andi. andis.  andc.  nand.  or. orc.  nor.  xor.  eqv.  extsb. extsh.  cntlzw.	cmp cmpi  cmpl cmpli	rlwimi.  rlwinm. rlwnm.	slw.  srw.  sraw. sawi.	crand crandc creqv crnand crnor cror crorc crxor  mcrf mcrxr mtrcf	tlbsx.	macchw.[o] macchws.[o] macchwsu.[o] macchwu.[o] machhw.[o] machhws.[o] machhwsu.[o] machhwu.[o] machlw.[o] machlws.[o] machlwsu.[o] machlwu.[o]  nmacchw.[o] nmacchws.[o] nmachhw.[o] nmachhws.[o] nmacchw.[o] nmacchws.[o]  mulchw. mulchwu. mulhww. mulhhwu. mullhw. mullhww.  dlmzb.

*Instruction Set* on page 243, provides detailed information on how each of these instructions updates the CR. To summarize, the CR can be accessed in any of the following ways:

- **mfcrr** reads the CR into a GPR. Note that this instruction does not *update* the CR and is therefore not listed in *Table 2-24*.
- Conditional branch instructions can designate a CR bit to be used as a branch condition. Note that these instructions do not *update* the CR and are therefore not listed in *Table 2-24*.
- **mtrcf** sets specified CR fields by writing to the CR from a GPR, under control of a mask field specified as part of the instruction.
- **mcrf** updates a specified CR field by copying another specified CR field into it.
- **mcrxr** copies certain bits of the XER into a specified CR field, and clears the corresponding XER bits.
- Integer compare instructions update a specified CR field.
- CR-logical instructions update a specified CR bit with the result of any one of eight logical operations on a specified pair of CR bits.

- Certain forms of various integer instructions (the “.” forms) implicitly update CR[CR0], as do certain forms of the auxiliary processor instructions implemented within the PPC440x5 core.
- Auxiliary processor instructions may in general update a specified CR field in an implementation-specified manner. In addition, if an auxiliary processor implements the floating-point operations specified by PowerPC Book-E, then those instructions update the CR in the manner defined by the architecture. See Book E: PowerPC Architecture Enhanced for Embedded Applications for details.

### *CR[CR0] Implicit Update By Integer Instructions*

Most of the CR-updating instructions listed in *Table 2-24* implicitly update the CR0 field. These are the various “dot-form” instructions, indicated by a “.” in the instruction mnemonic. Most of these instructions update CR[CR0] according to an arithmetic comparison of 0 with the 32-bit result which the instruction writes to the GPR file. That is, after performing the operation defined for the instruction, the 32-bit result which is written to the GPR file is compared to 0 using a signed comparison, independent of whether the actual operation being performed by the instruction is considered “signed” or not. For example, logical instructions such as **and.**, **or.**, and **nor.** update CR[CR0] according to this signed comparison to 0, even though the result of such a logical operation is not typically interpreted as a signed value. For each of these dot-form instructions, the individual bits in CR[CR0] are updated as follows:

CR[CR0] <sub>0</sub> — LT	Less than 0; set if the most-significant bit of the 32-bit result is 1.
CR[CR0] <sub>1</sub> — GT	Greater than 0; set if the 32-bit result is non-zero and the most-significant bit of the result is 0.
CR[CR0] <sub>2</sub> — EQ	Equal to 0; set if the 32-bit result is 0.
CR[CR0] <sub>3</sub> — SO	Summary overflow; a copy of XER[SO] at the completion of the instruction (including any XER[SO] update being performed the instruction itself).

Note that if an arithmetic overflow occurs, the “sign” of an instruction result indicated in CR[CR0] might not represent the “true” (infinitely precise) algebraic result of the instruction that set CR0. For example, if an **add.** instruction adds two large positive numbers and the magnitude of the result cannot be represented as a twos-complement number in a 32-bit register, an overflow occurs and CR[CR0]<sub>0</sub> is set, even though the infinitely precise result of the add is positive.

Similarly, adding the largest 32-bit twos-complement negative number (0x80000000) to itself results in an arithmetic overflow and 0x0000 0000 is recorded in the target register. CR[CR0]<sub>2</sub> is set, indicating a result of 0, but the infinitely precise result is negative.

CR[CR0]<sub>3</sub> is a copy of XER[SO] at the completion of the instruction, whether or not the instruction which is updating CR[CR0] is also updating XER[SO]. Note that if an instruction causes an arithmetic overflow but is not of the form which actually updates XER[SO], then the value placed in CR[CR0]<sub>3</sub> does not reflect the arithmetic overflow which occurred on the instruction (it is merely a copy of the value of XER[SO] which was already in the XER before the execution of the instruction updating CR[CR0]).

There are a few dot-form instructions which do not update CR[CR0] in the fashion described above. These instructions are: **stwcx.**, **tlbsx.**, and **d1mzb.** See the instruction descriptions in *Instruction Set* on page 243 for details on how these instructions update CR[CR0].

### CR Update By Integer Compare Instructions

Integer compare instructions update a specified CR field with the result of a comparison of two 32-bit numbers, the first of which is from a GPR and the second of which is either an immediate value or from another GPR. There are two types of integer compare instructions, *arithmetic* and *logical*, and they are distinguished by the interpretation given to the 32-bit numbers being compared. For *arithmetic* compares, the numbers are considered to be signed, whereas for *logical* compares, the numbers are considered to be unsigned. As an example, consider the comparison of 0 with 0xFFFFFFFF. In an *arithmetic* compare, 0 is larger; in a *logical* compare, 0xFFFFFFFF is larger.

A compare instruction can direct its result to any CR field. The BF field (bits 6:8) of the instruction specifies the CR field to be updated. After a compare, the specified CR field is interpreted as follows:

CR[(BF)] <sub>0</sub> — LT	The first operand is less than the second operand.
CR[(BF)] <sub>1</sub> — GT	The first operand is greater than the second operand.
CR[(BF)] <sub>2</sub> — EQ	The first operand is equal to the second operand.
CR[(BF)] <sub>3</sub> — SO	Summary overflow; a copy of XER[SO].

## 2.6 Integer Processing

Integer processing includes loading and storing data between memory and GPRs, as well as performing various operations on the values in GPRs and other registers (the categories of integer instructions are summarized in *Table 2-4* on page 56). The sections which follow describe the registers which are used for integer processing, and how they are updated by various instructions. In addition, *Condition Register (CR)* on page 66 provides more information on the CR updates caused by integer instructions. Finally, *Instruction Set* on page 243 also provides details on the various register updates performed by integer instructions.

### 2.6.1 General Purpose Registers (GPRs)

The PPC440x5 contains 32 GPRs. The contents of these registers can be transferred to and from memory using integer storage access instructions. Operations are performed on GPRs by most other instructions.

Access to the GPRs is non-privileged.




Figure 2-6. General Purpose Registers (R0-R31)



### 2.6.2 Integer Exception Register (XER)

The XER records overflow and carry indications from integer arithmetic and shift instructions. It also provides a byte count for string indexed integer storage access instructions (**lswx** and **stswx**). Note that the term *exception* in the name of this register does not refer to exceptions as they relate to interrupts, but rather to the *arithmetic* exceptions of carry and overflow.

Figure 2-7 illustrates the fields of the XER, while Table 2-25 and Table 2-26 list the instructions which update XER[SO,OV] and the XER[CA] fields, respectively. The sections which follow the figure and tables describe the fields of the XER in more detail.

Access to the XER is non-privileged.




Figure 2-7. Integer Exception Register (XER)

0	SO	Summary Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by <b>mtspr</b> or by integer or auxiliary processor instructions with the [o] option; can be <i>reset</i> by <b>mtspr</b> or by <b>mcrxr</b> .
1	OV	Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by <b>mtspr</b> or by integer or allocated instructions with the [o] option; can be <i>reset</i> by <b>mtspr</b> , by <b>mcrxr</b> , or by integer or allocated instructions with the [o] option.
2	CA	Carry 0 Carry has not occurred. 1 Carry has occurred.	Can be <i>set</i> by <b>mtspr</b> or by certain integer arithmetic and shift instructions; can be <i>reset</i> by <b>mtspr</b> , by <b>mcrxr</b> , or by certain integer arithmetic and shift instructions.
3:24		Reserved	
25:31	TBC	Transfer Byte Count	Used as a byte count by <b>lswx</b> and <b>stswx</b> ; written by <b>dlimzb[.]</b> and by <b>mtspr</b> .

Table 2-25. XER[SO,OV] Updating Instructions

Integer Arithmetic					Auxiliary Processor		Processor Control
Add	Subtract	Multiply	Divide	Negate	Multiply-Accumulate	Negative Multiply-Accumulate	Register Management
addo[.] addco[.] addeo[.] addmeo[.] addzeo[.]	subfo[.] subfco[.] subfeo[.] subfmeo[.] subfzeo[.]	mullwo[.]	divwo[.] divwuo[.]	nego[.]	macchw[.] macchwso[.] macchwsuo[.] macchwuo[.] machhwo[.] machhwsuo[.] machhwuo[.] maclhwo[.] maclhwsuo[.] maclhwuo[.]	nmacchw[.] nmacchwso[.] nmachhwo[.] nmachhwsuo[.] nmaclhwo[.] nmaclhwsuo[.]	mtspr mcrxr

Table 2-26. XER[CA] Updating Instructions

Integer Arithmetic		Integer Shift	Processor Control
Add	Subtract	Shift Right Algebraic	Register Management
addc[o][.] adde[o][.] addic[.] addme[o][.] addze[o][.]	subfc[o][.] subfe[o][.] subfc[.] subfme[o][.] subfze[o][.]	srav[.] srawi[.]	mtspr mcrxr

### 2.6.2.1 Summary Overflow (SO) Field

This field is set to 1 when an instruction is executed that causes XER[OV] to be set to 1, except for the case of **mtspr**(XER), which writes XER[SO,OV] with the values in (RS)<sub>0,1</sub>, respectively. Once set, XER[SO] is not reset until either an **mtspr**(XER) is executed with data that explicitly writes 0 to XER[SO], or until an **mcrxr** instruction is executed. The **mcrxr** instruction sets XER[SO] (as well as XER[OV,CA]) to 0 after copying all three fields into CR[CR0]<sub>0,2</sub> (and setting CR[CR0]<sub>3</sub> to 0).

Given this behavior, XER[SO] does not necessarily indicate that an overflow occurred on the most recent integer arithmetic operation, but rather that one occurred at some time subsequent to the last clearing of XER[SO] by **mtspr**(XER) or **mcrxr**.

XER[SO] is read (along with the rest of the XER) into a GPR by **mf spr**(XER). In addition, various integer instructions copy XER[SO] into CR[CR0]<sub>3</sub> (see *Condition Register (CR)* on page 66).

### 2.6.2.2 Overflow (OV) Field

This field is updated by certain integer arithmetic instructions to indicate whether the infinitely precise result of the operation can be represented in 32 bits. For those integer arithmetic instructions that update XER[OV] and produce *signed* results, XER[OV] = 1 if the result is greater than  $2^{31} - 1$  or less than  $-2^{31}$ ; otherwise,

XER[OV] = 0. For those integer arithmetic instructions that update XER[OV] and produce *unsigned* results (certain integer divide instructions and multiply-accumulate auxiliary processor instructions), XER[OV] = 1 if the result is greater than  $2^{32}-1$ ; otherwise, XER[OV] = 0. See the instruction descriptions in *Instruction Set* on page 243 for more details on the conditions under which the integer divide instructions set XER[OV] to 1.

The **mtspr**(XER) and **mcrxr** instructions also update XER[OV]. Specifically, **mcrxr** sets XER[OV] (and XER[SO,CA]) to 0 after copying all three fields into CR[CR0]<sub>0:2</sub> (and setting CR[CR0]<sub>3</sub> to 0), while **mtspr**(XER) writes XER[OV] with the value in (RS)<sub>1</sub>.

XER[OV] is read (along with the rest of the XER) into a GPR by **mfspir**(XER).

### 2.6.2.3 Carry (CA) Field

This field is updated by certain integer arithmetic instructions (the “carrying” and “extended” versions of add and subtract) to indicate whether or not there is a carry-out of the most-significant bit of the 32-bit result. XER[CA] = 1 indicates a carry. The integer shift right algebraic instructions update XER[CA] to indicate whether or not any 1-bits were shifted out of the least significant bit of the result, if the source operand was negative (see the instruction descriptions in *Instruction Set* on page 243 for more details).

The **mtspr**(XER) and **mcrxr** instructions also update XER[CA]. Specifically, **mcrxr** sets XER[CA] (as well as XER[SO,OV]) to 0 after copying all three fields into CR[CR0]<sub>0:2</sub> (and setting CR[CR0]<sub>3</sub> to 0), while **mtspr**(XER) writes XER[CA] with the value in (RS)<sub>2</sub>.

XER[CA] is read (along with the rest of the XER) into a GPR by **mfspir**(XER). In addition, the “extended” versions of the add and subtract integer arithmetic instructions use XER[CA] as a source operand for their arithmetic operations.

### Transfer Byte Count (TBC) Field

The TBC field is used by the string indexed integer storage access instructions (**lswx** and **stswx**) as a byte count. The TBC field is updated by the **dlmzb**[.] instruction with a value indicating the number of bytes up to and including the zero byte detected by the instruction (see the instruction description for **dlmzb** in *Instruction Set* on page 243 for more details). The TBC field is also written by **mtspr**(XER) with the value in (RS)<sub>25:31</sub>.

XER[TBC] is read (along with the rest of the XER) into a GPR by **mfspir**(XER).

## 2.7 Processor Control

The PPC440x5 core provides several registers for general processor control and status. These include:

- Machine State Register (MSR)
  - Controls interrupts and other processor functions
- Special Purpose Registers General (SPRGs)
  - SPRs for general purpose software use
- Processor Version Register (PVR)
  - Indicates the specific implementation of a processor
- Processor Identification Register (PIR)



Indicates the specific instance of a processor in a multi-processor system

- Core Configuration Register 0 (CCR0)

Controls specific processor functions, such as instruction prefetch

- Reset Configuration (RSTCFG)

Reports the values of certain fields of the TLB as supplied at reset

Except for the MSR, each of these registers is described in more detail in the following sections. The MSR is described in more detail in *Interrupts and Exceptions* on page 153.

### 2.7.1 Special Purpose Registers General (USPRG0, SPRG0–SPRG7)

USPRG0 and SPRG0–SPRG7 are provided for general purpose, system-dependent software use. One common system usage of these registers is as temporary storage locations. For example, a routine might save the contents of a GPR to an SPRG, and later restore the GPR from it. This is faster than a save/restore to a memory location. These registers are written using **mtspr** and read using **mfspir**.

Access to USPRG0 is non-privileged for both read and write.

Access to SPRG4–SPRG7 is non-privileged for read but privileged for write, and hence different SPR numbers are used for reading than for writing.

Access to SPRG0–SPRG3 is privileged for both read and write.




Figure 2-8. Special Purpose Registers General (USPRG0, SPRG0–SPRG7)

0:31	General data	Software value; no hardware usage.
------	--------------	------------------------------------

### 2.7.2 Processor Version Register (PVR)

The PVR is a read-only register typically used to identify a specific processor core and chip implementation. Software can read the PVR to determine processor core and chip hardware features. The PVR can be read into a GPR using **mfspir**.

Refer to PowerPC 440x5 Embedded Processor Data Sheet for the PVR value.

Access to the PVR is privileged.




Figure 2-9. Processor Version Register (PVR)

0:11	OWN	Owner Identifier	Identifies the owner of a core.
12:31	PVN	Processor Version Number	Implementation-specific value identifying the specific version and use of a processor core within a chip.

### 2.7.3 Processor Identification Register (PIR)

The PIR is a read-only register that uniquely identifies a specific instance of a processor core, within a multi-processor configuration, enabling software to determine exactly which processor it is running on. This capability is important for operating system software within multiprocessor configurations. The PIR can be read into a GPR using **mfspr**.

Because the PPC440x5 is a uniprocessor, PIR[PIN] = 0b0000.

Access to the PIR is privileged.



Figure 2-10. Processor Identification Register (PIR)

0:27		Reserved
28:31	PIN	Processor Identification Number (PIN)

### 2.7.4 Core Configuration Register 0 (CCR0)

The CCR0 controls a number of special chip functions, including data cache and auxiliary processor operation, speculative instruction fetching, trace, and the operation of the cache block touch instructions. The CCR0 is written from a GPR using **mtspr**, and can be read into a GPR using **mfspr**. Figure 2-11 illustrates the fields of the CCR0, and gives a brief description of their functions. A cross reference after the bit-field description indicates the section of this document which describes each field in more detail.

Access to the CCR0 is privileged.




Figure 2-11. Core Configuration Register 0 (CCR0)

0		Reserved	
1	PRE	Parity Recoverability Enable 0 Semi-recoverable parity mode enabled for data cache 1 Fully recoverable parity mode enabled for data cache	Must be set to 1 to guarantee full recoverability from MMU and data cache parity errors.
2:3		Reserved	
4	CRPE	Cache Read Parity Enable 0 Disable parity information reads 1 Enable parity information reads	When enabled, execution of <b>icread</b> , <b>dcread</b> , or <b>tlbre</b> loads parity information into the ICDBTRH, DCDBTRL, or target GPR, respectively.
5:9		Reserved	
10	DSTG	Disable Store Gathering 0 Enabled; stores to contiguous addresses may be gathered into a single transfer 1 Disabled; all stores to memory will be performed independently	See <i>Store Gathering</i> on page 116.
11	DAPUIB	Disable APU Instruction Broadcast 0 Enabled. 1 Disabled; instructions not broadcast to APU for decoding	This mechanism is provided as a means of reducing power consumption when an auxilliary processor is not attached and/or is not being used. See <i>Initialization</i> on page 83.
12:15		Reserved	
16	DTB	Disable Trace Broadcast 0 Enabled. 1 Disabled; no trace information is broadcast.	This mechanism is provided as a means of reducing power consumption when instruction tracing is not needed. See <i>Initialization</i> on page 83.
17	GICBT	Guaranteed Instruction Cache Block Touch 0 <b>icbt</b> may be abandoned without having filled cache line if instruction pipeline stalls. 1 <b>icbt</b> is guaranteed to fill cache line even if instruction pipeline stalls.	See <i>icbt Operation</i> on page 108.
18	GDCBT	Guaranteed Data Cache Block Touch 0 <b>dcbt/dcbtst</b> may be abandoned without having filled cache line if load/store pipeline stalls. 1 <b>dcbt/dcbtst</b> are guaranteed to fill cache line even if load/store pipeline stalls.	See <i>Data Cache Control and Debug</i> on page 121.
19:22		Reserved	
23	FLSTA	Force Load/Store Alignment 0 No Alignment exception on integer storage access instructions, regardless of alignment 1 An alignment exception occurs on integer storage access instructions if data address is not on an operand boundary.	See <i>Load and Store Alignment</i> on page 114.
24:27		Reserved	

28:29	ICSLC	Instruction Cache Speculative Line Count	Number of additional lines (0–3) to fill on instruction fetch miss. See <i>Speculative Prefetch Mechanism</i> on page 102.
30:31	ICSLT	Instruction Cache Speculative Line Threshold	Number of doublewords that must have already been filled in order that the current speculative line fill is <i>not</i> abandoned on a redirection of the instruction stream. See <i>Speculative Prefetch Mechanism</i> on page 102.

### 2.7.5 Core Configuration Register 1 (CCR1)

Bits 0:19 of CCR1 can cause all possible parity error exceptions to verify correct machine check exception handler operation. Other CCR1 bits can force a full-line data cache flush, or select a CPU timer clock input other than CPUClock. The CCR1 is written from a GPR using **mtspr**, and can be read into a GPR using **mfspr**. *Figure 2-12* illustrates the fields of the CCR1, and gives a brief description of their functions.

Access to the CCR1 is privileged.




Figure 2-12. Core Configuration Register 1 (CCR1)

0:7	ICDPEI	Instruction Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded when the instruction cache is filled. Each of the 8 bits corresponds to one of the instruction words in the line.
8:9	ICTPEI	Instruction Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the instruction cache.
10:11	DCTPEI	Data Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the data cache.
12	DCDPEI	Data Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the data field in the data cache.
13	DCUPEI	Data Cache U-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bit recorded for the U fields in the data cache.
14	DCMPEI	Data Cache Modified-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the modified (dirty) field in the data cache.
15	FCOM	Force Cache Operation Miss 0 normal operation 1 cache ops appear to miss the cache	Force <b>icbt</b> , <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>dcbi</b> , and <b>dcbz</b> to appear to miss the caches. The intended use is with <b>icbt</b> and <b>dcbt</b> only, which will fill a duplicate line and allow testing of multi-hit parity errors. See <i>Section 4.2.4.7 Simulating Instruction Cache Parity Errors for Software Testing</i> on page 111 and <i>Figure 4.3.3.7</i> on page 126.

16:19	MMUPEI	Memory Management Unit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the MMU.
20	FFF	Force Full-line Flush 0 flush only as much data as necessary. 1 always flush entire cache lines	When flushing 32-byte (8-word) lines from the data cache, normal operation is to write nothing, a double word, quad word, or the entire 8-word block to the memory as required by the dirty bits. This bit ensures that none or all dirty bits are set so that either nothing or the entire 8-word block is written to memory when flushing a line from the data cache. Refer to <i>Section 4.3.1.4 Line Flush Operations</i> on page 117.
21:23		Reserved	
24	TCS	Timer Clock Select 0 CPU timer advances by one at each rising edge of the CPU input clock (CPMC440CLOCK). 1 CPU timer advances by one for each rising edge of the CPU timer clock (CPMC440TIMERCLOCK).	When TCS = 1, CPU timer clock input can toggle at up to half of the CPU clock frequency.
25:31		Reserved	

## 2.7.6 Reset Configuration (RSTCFG)

The read-only RSTCFG register reports the values of certain fields of TLB as supplied at reset.

Access to RSTCFG is privileged.




Figure 2-13. Reset Configuration

0:15		Reserved	
16	U0	U0 Storage Attribute 0 U0 storage attribute is disabled 1 U0 storage attribute is enabled	See <i>Table 5-1</i> on page 131.
17	U1	U1 Storage Attribute 0 U1 storage attribute is disabled 1 U1 storage attribute is enabled	See <i>Table 5-1</i> on page 131.
18	U2	U2 Storage Attribute 0 U2 storage attribute is disabled 1 U2 storage attribute is enabled	See <i>Table 5-1</i> on page 131.
19	U3	U3 Storage Attribute 0 U3 storage attribute is disabled 1 U3 storage attribute is enabled	See <i>Table 5-1</i> on page 131.
20:23		Reserved	

24	E	E Storage Attribute 0 Accesses to the page are big endian. 1 Accesses to the page are little endian.	
25:27		Reserved	
28:31	ERPn	Extended Real Page Number	This TLB field is prepended to the translated address to form a 36-bit real address. See <i>Table 5.4 Address Translation</i> on page 136 and <i>Table 5-3 Page Size and Real Address Formation</i> on page 137.

## 2.8 User and Supervisor Modes

PowerPC Book-E architecture defines two operating “states” or “modes,” supervisor (privileged), and user (non-privileged). Which mode the processor is operating in is controlled by MSR[PR]. When MSR[PR] is 0, the processor is in supervisor mode, and can execute all instructions and access all registers, including privileged ones. When MSR[PR] is 1, the processor is in user mode, and can only execute non-privileged instructions and access non-privileged registers. An attempt to execute a privileged instruction or to access a privileged register while in user mode causes a Privileged Instruction exception type Program interrupt to occur.

Note that the name “PR” for the MSR field refers to an historical alternative name for user mode, which is “problem state.” Hence the value 1 in the field indicates “problem state,” and not “privileged” as one might expect.

### 2.8.1 Privileged Instructions

The following instructions are privileged and cannot be executed in user mode:

Table 2-27. Privileged Instructions

dcbi	
dccci	
dcread	
iccci	
icread	
mfocr	
mfspr	For any SPR Number with SPRN <sub>5</sub> = 1. See <i>Privileged SPRs</i> on page 79.
mtocr	
mtmsr	
mtspr	For any SPR Number with SPRN <sub>5</sub> = 1. See <i>Privileged SPRs</i> on page 79.
rfci	
rfi	
rfmci	
tlbre	
tlbsx	
tlbsync	

Table 2-27. Privileged Instructions (continued)

tlbwe	
wrtee	
wrteei	

## 2.8.2 Privileged SPRs

Most SPRs are privileged. The only defined non-privileged SPRs are the LR, CTR, XER, USPRG0, SPRG4–7 (read access only), TBU (read access only), and TBL (read access only). The PPC440x5 core also treats all SPR numbers with a 1 in bit 5 of the SPRN field as privileged, whether the particular SPR number is defined or not. Thus the core causes a Privileged Instruction exception type Program interrupt on any attempt to access such an SPR number while in user mode. In addition, the core causes an Illegal Instruction exception type Program interrupt on any attempt to access while in user mode an undefined SPR number with a 0 in SPRN<sub>5</sub>. On the other hand, the result of attempting to access an undefined SPR number in supervisor mode is undefined, regardless of the value in SPRN<sub>5</sub>.

## 2.9 Speculative Accesses

The PowerPC Book-E Architecture permits implementations to perform speculative accesses to memory, either for instruction fetching, or for data loads. A speculative access is defined as any access that is not required by the sequential execution model (SEM).

For example, the PPC440x5 speculatively prefetches instructions down the predicted path of a conditional branch; if the branch is later determined to not go in the predicted direction, the fetching of the instructions from the predicted path is not required by the SEM and thus is speculative. Similarly, the PPC440x5 executes load instructions out-of-order, and may read data from memory for a load instruction that is past an undetermined branch.

Sometimes speculative accesses are inappropriate, however. For example, attempting to access data at addresses to which I/O devices are mapped can cause problems. If the I/O device is a serial port, reading it speculatively could cause data to be lost.

The architecture provides two mechanisms for protecting against errant accesses to such “non-well-behaved” memory addresses. The first is the guarded (G) storage attribute, and protects against speculative data accesses. The second is the execute permission mechanism, and protects against speculative instruction fetches. Both of these mechanisms are described in *Memory Management* on page 129

## 2.10 Synchronization

The PPC440x5 supports the synchronization operations of the PowerPC Book-E architecture. There are three kinds of synchronization defined by the architecture, each of which is described in the following sections.

### 2.10.1 Context Synchronization

The context of a program is the environment in which the program executes. For example, the mode (user or supervisor) is part of the context, as are the address translation space and storage attributes of the memory pages being accessed by the program. Context is controlled by the contents of certain registers and other resources, such as the MSR and the translation lookaside buffer (TLB).

Under certain circumstances, it is necessary for the hardware or software to force the synchronization of a program's context. Context synchronizing operations include all interrupts except Machine Check, as well as the **isync**, **sc**, **rfi**, **rfci**, and **rfmci** instructions. Context synchronizing operations satisfy the following requirements:

1. The operation is not initiated until all instructions preceding the operation have completed to the point at which they have reported any and all exceptions that they will cause.
2. All instructions *preceding* the operation must complete in the context in which they were initiated. That is, they must not be affected by any context changes caused by the context synchronizing operation, or any instructions *after* the context synchronizing operation.
3. If the operation is the **sc** instruction (which causes a System Call interrupt) or is itself an interrupt, then the operation is not initiated until no higher priority interrupt is pending (see *Interrupts and Exceptions* on page 153).
4. All instructions that *follow* the operation must be re-fetched and executed in the context that is established by the completion of the context synchronizing operation and all of the instructions which *preceded* it.

Note that context synchronizing operations do not force the completion of storage accesses, nor do they enforce any ordering amongst accesses before and/or after the context synchronizing operation. If such behavior is required, then a storage synchronizing instruction must be used (see *Storage Ordering and Synchronization* on page 81).

Also note that architecturally Machine Check interrupts are not context synchronizing. Therefore, an instruction that *precedes* a context synchronizing operation can cause a Machine Check interrupt *after* the context synchronizing operation occurs and additional instructions have completed. For the PPC440x5 core, this can only occur with Data Machine Check exceptions, and not Instruction Machine Check exceptions.

The following scenarios use pseudocode examples to illustrate the effects of context synchronization. Subsequent text explains how software can further guarantee "storage ordering."

1. Consider the following self-modifying code instruction sequence:

```
stw XYZ   Store to caching inhibited address XYZ
isync
XYZ       fetch and execute the instruction at address XYZ
```

In this sequence, the **isync** instruction does not guarantee that the XYZ instruction is fetched after the store has occurred to memory. There is no guarantee which XYZ instruction will execute; either the old version or the new (stored) version might.

2. Now consider the required self-modifying code sequence:

```
stw       Write new instruction to data cache
dcbst     Push the new instruction from the data cache to memory
msync     Guarantee that dcbst completes before subsequent instructions begin
icbi     invalidate old copy of instruction in instruction cache
```



**msync** Guarantee that **icbi** completes before subsequent instructions begin  
**isync** force context synchronization, discard ed instructions and re-fetch, fetch of

PVR	OWN	System-dependent	PVR[OWN] value (after reset and otherwise) is specified by core input signals
	PVN	System-dependent	PVR[PVN] value (after reset and otherwise) is specified by core input signals

stored instruction guaranteed to get new value

3. This final example illustrates the use of **isync** with context changes to the debug facilities

**mtdbcr0** Enable the instruction address compare (IAC) debug event  
**isync** Wait for the new Debug Control Register 0 (DBCR0) context to be established  
**XYZ** This instruction is at the IAC address; an **isync** is necessary to guarantee that the IAC event is recognized on the execution of this instruction; without the **isync**, the XYZ instruction may be prefetched and dispatched to execution before recognizing that the IAC event has been enabled.

### 2.10.2 Execution Synchronization

Execution synchronization is a subset of context synchronization. An execution synchronizing operation satisfies the first two requirements of context synchronizing operations, but not the latter two. That is, execution synchronizing operations guarantee that preceding instructions execute in the “old” context, but do not guarantee that subsequent instructions operate in the “new” context. An example of a scenario requiring execution synchronization would be just before the execution of a TLB-updating instructions (such as **tlbwe**). An execution synchronizing instruction should be executed to guarantee that all preceding storage access instructions have performed their address translations before executing **tlbwe** to invalidate an entry which might be used by those preceding instructions.

There are four execution synchronizing instructions: **mtmsr**, **wrtee**, **wrteei**, and **msync**. Of course, all context synchronizing instruction are also implicitly execution synchronizing, since context synchronization is a superset of execution synchronization.

Note that PowerPC Book-E imposes additional requirements on updates to MSR[EE] (the external interrupt enable bit). Specifically, if a **mtmsr**, **wrtee**, or **wrteei** instruction sets MSR[EE] = 1, and an External Input, Decrementer, or Fixed Interval Timer exception is pending, the interrupt must be taken before the instruction that follows the MSR[EE]-updating is executed. In this sense, these MSR[EE]-updating instructions can be thought of as being context synchronizing with respect to the MSR[EE] bit, in that it guarantees that subsequent instructions execute (or are prevented from executing and an interrupt taken) according to the new context of MSR[EE].

### 2.10.3 Storage Ordering and Synchronization

Storage synchronization enforces ordering between storage access instructions executed by the PPC440x5 core. There are two storage synchronizing instructions: **msync** and **mbar**. PowerPC Book-E architecture defines different ordering requirements for these two instructions, but the PPC440x5 core implements them in an identical fashion. Architecturally, **msync** is the “stronger” of the two, and is also execution synchronizing, whereas **mbar** is not.

**mbar** acts as a “barrier” between all storage access instructions executed before the **mbar** and all those executed after the **mbar**. That is, **mbar** ensures that all of the storage accesses initiated by instructions before the **mbar** are performed with respect to the memory subsystem before any of the accesses initiated by

instructions after the **mbar**. However, **mbar** does not prevent subsequent instructions from executing (nor even from completing) before the completion of the storage accesses initiated by instructions before the **mbar**.

**msync**, on the other hand, does guarantee that all preceding storage accesses have actually been performed with respect to the memory subsystem before the execution of any instruction after the **msync**. Note that this requirement goes beyond the requirements of mere execution synchronization, in that execution synchronization doesn't require the completion of preceding storage accesses.

The following two examples illustrate the distinctive use of **mbar** vs. **msync**.

stw	Store data to an I/O device
msync	Wait for store to actually complete
mtdcr	Reconfigure the I/O device

In this example, the **mtdcr** is reconfiguring the I/O device in a manner which would cause the preceding store instruction to fail, were the **mtdcr** to change the device before the completion of the store. Since **mtdcr** is not a storage access instruction, the use of **mbar** instead of **msync** would not guarantee that the store is performed before letting the **mtdcr** reconfigure the device. It only guarantees that subsequent storage accesses are not performed to memory or any device before the earlier store.

Now consider this next example:

stb X	Store data to an I/O device at address X, causing a status bit at address Y to be reset
mbar	Guarantee preceding store is performed to the device before any subsequent storage accesses are performed
lbz Y	Load status from the I/O device at address Y

Here, **mbar** is appropriate instead of **msync**, because all that is required is that the store to the I/O device happens before the load does, but not that other instructions subsequent to the **mbar** won't get executed before the store.

## 3. Initialization

This chapter describes the initial state of the PPC440x5 core after a hardware reset, and contains a description of the initialization software required to complete initialization so that the PPC440x5 core can begin executing application code. Initialization of other on-chip and/or off-chip system components may also be needed, in addition to the processor core initialization described in this chapter.

### 3.1 PPC440x5 Core State After Reset

In general, the contents of registers and other facilities within the PPC440x5 core are undefined after a hardware reset. Reset is defined to initialize only the minimal resources required such that instructions can be fetched and executed from the initial program memory page, and so that repeatable, deterministic behavior can be guaranteed provided that the proper software initialization sequence is followed. System software must fully configure the rest of the PPC440x5 core resources, as well as the other facilities within the chip and/or system.

The following list summarizes the requirements of the Book-E Enhanced PowerPC Architecture with regards to the processor state after reset, prior to any additional initialization by software.

- All fields of the MSR are set to 0, disabling all asynchronous interrupts, placing the processor in supervisor mode, and specifying that instruction and data accesses are to the system (as opposed to application) address space.
- DBCR0[RST] is set to 0, thereby ending any previous software-initiated reset operation.
- DBSR[MRR] records the type of the just ended reset operation (core, chip, or system; see *Reset Types* on page 87).
- TCR[WRC] is set to 0, thereby disabling the Watchdog timer reset operation.
- TSR[WRS] records the type of the just ended reset operation, if the reset was initiated by the Watchdog Timer (otherwise this field is unchanged from its pre-reset value).
- The PVR is defined, after reset and otherwise, to contain a value that indicates the specific processor implementation.
- The program counter (PC) is set to 0xFFFFF0FC, the effective address (EA) of the last word of the address space.

The memory management resources are set to values such that the processor is able to successfully fetch and execute instructions and read (but not write) data within the 4KB program memory page located at the end of the 32-bit effective address space. Exactly how this is accomplished is implementation-dependent. For example, it may or may not be the case that a TLB entry is established in a manner which is visible to software using the TLB management instructions. Regardless of how the implementation enables access to the initial program memory page, instruction execution starts at the effective address of 0xFFFFF0FC, the last word of the effective address space. The instruction at this address must be an unconditional branch backwards to the start of the initialization sequence, which must lie somewhere within the initial 4KB program memory page. The real address to which the initial effective address will be translated is also implementation- or system-dependent, as are the various storage attributes of the initial program memory page such as the caching inhibited and endian attributes.

**Note:** In the PPC440x5 core, a single entry is established in the instruction shadow TLB (ITLB) and data shadow TLB (DTLB) at reset with the properties described in Table 3-1. It is required that initialization software insert an entry into the UTLB to cover this same memory region before performing any context synchro-

nizing operation (including causing any exceptions which would lead to an interrupt), since a context synchronizing operation will invalidate the shadow TLB entries.

Initialization software should consider all other resources within the PPC440x5 core to be undefined after reset, in order for the initialization sequence to be compatible with other PowerPC implementations. There are, however, additional core resources which are initialized by reset, in order to guarantee correct and deterministic operation of the processor during the initialization sequence. *Table 3-1* shows the reset state of all PPC440x5 core resources which are defined to be initialized by reset. While certain other register fields and other facilities within the PPC440x5 core may be affected by reset, this is not an architectural nor hardware requirement, and software must treat those resources as undefined. Likewise, even those resources which are included in *Table 3-1* but which are not identified in the previous list as being architecturally required, should be treated as undefined by the initialization software.

During chip initialization, some chip control registers must be initialized to ensure proper chip operation. Peripheral devices can also be initialized as appropriate for the system design.

*Table 3-1. Reset Values of Registers and Other PPC440x5 Facilities*

Resource	Field	Reset Value	Comment
CCR0	DAPUIB	0	Enable broadcast of instruction data to auxiliary processor interface
	DTB	0	Enable broadcast of trace information
CCR1	ICDPEI	0	Disable Parity Error Insertion (enabled only for s/w testing)
	ICTPEI	0	
	DCTPEI	0	
	DCDPEI	0	
	DCUPEI	0	
	DCMPEI	0	
	FCOM	0	Do not force cache ops to miss.
	MMUPEI	0	Disable Parity Error Insertion (enabled only for s/w testing)
	FFF	0	Flush only as much data from dirty lines as needed.
DBCR0	EDM	0	External Debug mode disabled
	RST	0b00	Software-initiated debug reset disabled
	ICMP	0	Instruction completion debug events disabled
	BRT	0	Branch taken debug events disabled
	IAC1	0	Instruction Address Compare 1 (IAC1) debug events disabled
	IAC2	0	IAC2 debug events disabled
	IAC3	0	IAC3 debug events disabled
	IAC4	0	IAC4 debug events disabled

Table 3-1. Reset Values of Registers and Other PPC440x5 Facilities

Resource	Field	Reset Value	Comment
DBSR	UDE	0	Unconditional debug event has not occurred
	MRR	Reset-dependent	Indicates most recent type of reset as follows: 00 No reset has occurred since this field last cleared by software 01 Core reset 10 Chip reset 11 System reset
	ICMP	0	Instruction completion debug event has not occurred
	BRT	0	Branch taken debug event has not occurred
	IRPT	0	Interrupt debug event has not occurred
	TRAP	0	Trap debug event has not occurred
	IAC1	0	IAC1 debug event has not occurred
	IAC2	0	IAC2 debug event has not occurred
	IAC3	0	IAC3 debug event has not occurred
	IAC4	0	IAC4 debug event has not occurred
	DAC1R	0	Data address compare 1 (DAC1) read debug event has not occurred
	DAC1W	0	DAC1 write debug event has not occurred
	DAC2R	0	DAC2 read debug event has not occurred
	DAC2W	0	DAC2 write debug event has not occurred
	RET	0	Return debug event has not occurred
ESR	MCI	0	Synchronous Instruction Machine Check exception has not occurred
MCSR	MCS	0	Asynchronous Instruction Machine Check exception has not occurred
MSR	WE	0	Wait state disabled
	CE	0	Asynchronous critical interrupts disabled
	EE	0	Asynchronous non-critical interrupts disabled
	PR	0	Processor in supervisor mode
	FP	0	Floating-point Unavailable interrupts disabledStorage
	ME	0	Machine Check interrupts disabled
	FE0	0	Floating-point Enabled interrupts disabled
	DWE	0	Debug Wait mode disabled
	DE	0	Debug interrupts disabled
	FE1	0	Floating-point Enabled interrupts disabled
	IS	0	Instruction fetch access is to system-level virtual address space
	DS	0	Data access is to system level virtual address space
PC		0xFFFFF0FC	Initial reset instruction fetched from last word of effective address space
PVR	OWN	System-dependent	PVR[OWN] value (after reset and otherwise) is specified by core input signals
	PVN	System-dependent	PVR[PVN] value (after reset and otherwise) is specified by core input signals

Table 3-1. Reset Values of Registers and Other PPC440x5 Facilities

Resource	Field	Reset Value	Comment
RSTCFG	U0	System-dependent	All RSTCFG fields are specified by core input signals
	U1	System-dependent	
	U2	System-dependent	
	U3	System-dependent	
	E	System-dependent	
	EPRN	System-dependent	
TCR	WRC	0b00	Watchdog Timer reset disabled
TLBentry <sup>1</sup>	EPN <sub>0:19</sub>	0xFFFFF	Match EA of initial reset instruction (EPN <sub>20:21</sub> are undefined, as they are not compared to the EA because the page size is 4KB).
	V	1	Translation table entry for the initial program memory page is valid.
	TS	0	Initial program memory page is in system-level virtual address space.
	SIZE	0b0001	Initial program memory page size is 4KB.
	TID	0x00	Initial program memory page is globally shared; no match required against PID register.
	RPN <sub>0:21</sub>	0xFFFF    0b00	Initial program memory page mapped effective=real.
	ERPN	System-dependent	Extended real page number of the initial program memory page is specified by core input signals.
	U0–U3	System-dependent	Reset value of user-definable storage attributes are specified by core input signals
	W	0	Write-through storage attribute disabled.
	I	1	Caching inhibited storage attribute enabled.
	M	0	Memory coherent storage attribute disabled.
	G	1	Guarded storage attribute enabled.
	E	System-dependent	Reset value of endian storage attribute is specified by a core input signal.
	SX	1	Supervisor mode execution access enabled.
	SW	0	Supervisor mode write access disabled.
SR	1	Supervisor mode read access enabled.	
TSR	WRS	Copy of TCR[WRC]	If reset caused by Watchdog Timer
		Unchanged	If reset not caused by Watchdog Timer
		Undefined	After power-up

**Note:** "TLBentry" refers to an entry in the shadow instruction and data TLB arrays that is automatically configured by the PPC440x5 core to enable fetching and reading (but not writing) from the initial program memory page. This entry is not architecturally visible to software, and is invalidated upon any context synchronizing operation. Software must initialize a corresponding entry in the main unified TLB array before executing any operation which could lead to a context synchronization. See *Initialization Software Requirements* on page 87 for more information.

## 3.2 Reset Types

The PPC440x5 core supports three types of reset: core, chip, and system. The type of reset is indicated by a set of core input signals. For each type of reset, the core resources are initialized as indicated in *Table 3-1* on page 84. Core reset is intended to reset the PPC440x5 core without necessarily resetting the rest of the on-chip logic. The chip reset operation is intended to reset the entire chip, but off-chip hardware in the system is not informed of the reset operation. System reset is intended to reset the entire chip, and also to signal the rest of the off-chip system that the chip is being reset.

## 3.3 Reset Sources

A reset operation can be initiated on the PPC440x5 core through the use of any of four separate mechanisms. The first is a set of three input signals to the core, one for each of the three reset types. These signals can be asserted asynchronously by hardware outside the core to initiate a reset operation. The second reset source is the TCR[WRC] field, which can be setup by software to initiate a reset operation upon certain Watchdog Timer expiration events. The third reset source is the DBCR0[RST] field, which can be written by software to immediately initiate a reset operation. The fourth reset source is the JTAG interface, which can be used by a JTAG-attached debug tool to initiate a reset operation asynchronously to program execution on the PPC440x5 core.

## 3.4 Initialization Software Requirements

After a reset operation occurs, the PPC440x5 core is initialized to a minimum configuration to enable the fetching and execution of the software initialization code, and to guarantee deterministic behavior of the core during the execution of this code. Initialization software is necessary to complete the configuration of the processor core and the rest of the on-chip and off-chip system.

The system must provide non-volatile memory (or memory initialized by some mechanism other than the PPC440x5 core) at the real address corresponding to effective address 0xFFFFF0FC, and at the rest of the initial program memory page. The instruction at the initial address must be an unconditional branch backwards to the beginning of the initialization software sequence.

The initialization software functions described in this section perform the configuration tasks required to prepare the PPC440x5 core to boot an operating system and subsequently execute an application program.

The initialization software must also perform functions associated with hardware resources that are outside the PPC440x5 core, and hence that are beyond the scope of this manual. This section makes reference to some of these functions, but their full scope is described in the user's manual for the specific chip and/or system implementation.

Initialization software should perform the following tasks in order to fully configure the PPC440x5 core. For more information on the various functions referenced in the initialization sequence, see the corresponding chapters of this document.

1. Branch backwards from effective address 0xFFFFF0FC to the start of the initialization sequence
2. Invalidate the instruction cache (**iccci**)

3. Invalidate the data cache (**dccci**)4. Synchronize memory accesses (**msync**)

This step forces any data PLB operations that may have been in progress prior to the reset operation to complete, thereby allowing subsequent data accesses to be initiated and completed properly.

## 5. Clear DBCR0 register (disable all debug events)

Although the PPC440x5 core is defined to reset some of the debug event enables during the reset operation (as specified in *Table 3-1* on page 84), this is not required by the architecture and hence the initialization software should not assume this behavior. Software should disable all debug events in order to prevent non-deterministic behavior on the trace interface to the core.

## 6. Clear DBSR register (initialize all debug event status)

Although the PPC440x5 core is defined to reset the DBSR debug event status bits during the reset operation (as specified in *Table 3-1* on page 84), this is not required by the architecture and hence the initialization software should not assume this behavior. Software should clear all such status in order to prevent non-deterministic behavior on the JTAG interface to the core.

## 7. Initialize CCR0 register

1. Enable/disable broadcast of instructions to auxiliary processor (save power if no AP attached)
2. Enable/disable broadcast of trace information (save power if not tracing)
3. Enable/configure or disable speculative instruction cache line prefetching
4. Specify behavior for **icbt** and **dcbt/dcbtst** instructions
5. Enable/disable gathering of separate store accesses
6. Enable/disable hardware support for misaligned data accesses
7. Enable/disable parity error recoverability (recoverability lowers load/store performance marginally.)
8. Enable/disable cache read of parity bits depending on s/w compatibility requirements

## 8. Initialize CCR1 register

1. enable/disable full-line flushes as desired.
2. disable force cache-op miss (FCOM) and various parity error insertion (xxxPEI).
3. Users may wish to initialize CCR1[TCS] here, or in the timer facilities section.

## 9. Configure instruction and data cache regions

These steps must be performed prior to enabling the caches by setting the caching inhibited storage attribute of the corresponding TLB entry to 0.

1. Clear the instruction and data cache normal victim index registers (INV0–INV3, DNV0–DNV3)
2. Clear the instruction and data cache transient victim index registers (ITV0–ITV3, DTV0–DTV3)
3. Set the instruction and data cache victim limit registers (IVLIM and DVLIM) according to the desired size of the normal, locked, and transient regions of each cache

## 10. Setup TLB entry to cover initial program memory page

Since the PPC440x5 core only initializes an architecturally-invisible shadow TLB entry during the reset operation, and since all shadow TLB entries are invalidated upon any context synchronization, special care must be taken during the initialization sequence to prevent any such context synchronizing opera-



tions (such as interrupts and the **isync** instruction) until after this step is completed, and an architected TLB entry has been established in the TLB. Particular care should be taken to avoid store operations, since write permission is disabled upon reset, and an attempt to execute any store operation would result in a Data Storage interrupt, thereby invalidating the shadow TLB entry.

1. Initialize MMUCR

- Specify TID field to be written to TLB entries
- Specify TS field to be used for TLB searches
- Specify store miss allocation behavior
- Enable/disable transient cache mechanism
- Enable/disable cache locking exceptions

2. Write TLB entry for initial program memory page

- Specify EPN, RPN, ERPN, and SIZE as appropriate for system
- Set valid bit
- Specify TID = 0 (disable comparison to PID) or else initialize PID register to matching value
- Specify TS = 0 (system address space) or else MSR[IS,DS] must be set to correspond to TS=1
- Specify storage attributes (W, I, M, G, E, U0–U3) as appropriate for system
- Enable supervisor mode fetch, read, and write access (SX, SR, SW)

3. Initialize PID register to match TID field of TLB entry (unless using TID = 0)

4. Setup for subsequent MSR[IS,DS] initialization to correspond to TS field of TLB entry

Only necessary if TS field of TLB entry being set to 1 (MSR[IS,DS] already reset to 0)

- Write new MSR value into SRR1
- Write address from which to continue execution into SRR0

5. Setup for subsequent change in instruction fetch address

Only necessary if EPN field of TLB entry changed from the initial value ( $EPN_{0:19} \neq 0xFFFFF$ )

- Write initial/new MSR value into SRR1
- Write address from which to continue execution into SRR0

6. Fully initialize the TLB (tlbwe to all three words of each TLB entry; tlbre to TLB entries that are not fully initialized may result in parity exceptions).

7. Context synchronize to invalidate shadow TLB contents and cause new TLB contents to take effect

- Use **isync** if not changing MSR contents and not changing the effective address of the rest of the initialization sequence
- Use **rfi** if changing MSR to match new TS field of TLB entry (SRR1 will be copied into MSR, and program execution will resume at value in SRR0)
- Use **rfi** if changing next instruction fetch address to correspond to new EPN field of TLB entry (SRR1 will be copied into MSR, and program execution will resume at value in SRR0)

Instruction and data caches will now begin to be used, if the corresponding TLB entry has been setup with the caching inhibited storage attribute set to 0. Initialization software can now branch outside of

the initial 4KB memory region as controlled by the address and size of the new TLB entry and/or any other TLB entries which have been setup.

#### 11. Initialize interrupt resources

1. Initialize IVPR to specify high-order address of the interrupt handling routines

Make sure that the corresponding address region is covered by a TLB entry (or entries)

2. Initialize IVOR0–IVOR15 registers (individual interrupt vector addresses)

Make sure that the corresponding addresses are covered by a TLB entry (or entries)

Because the low order four bits of IVOR0–IVOR15 are reserved, the values written to those bits are ignored when the registers are written, and are read as zero when the registers are used. Therefore, all interrupt vector offsets are implicitly aligned on quadword boundaries. Software must take care to assure that all interrupt handlers are quadword-aligned.

3. Setup corresponding memory contents with the interrupt handling routines
4. Synchronize any program memory changes as required. (See *Self-Modifying Code* on page 103 for more information on the instruction sequence necessary to synchronize changes to program memory prior to executing the new instructions.)

#### 12. Configure debug facilities as desired

1. Write DBCR1 and DBCR2 to specify IAC and DAC event conditions
2. Clear DBSR to initialize IAC auto-toggle status
3. Initialize IAC1–IAC4, DAC1–DAC2, DVC1–DVC2 registers to desired values
4. Write MSR[DWE] to enable Debug Wait mode (if desired)
5. Write DBCR0 to enable desired debug mode(s) and event(s)
6. Context synchronize to establish new debug facility context (**isync**)

#### 13. Configure timer facilities as desired

1. Write DEC to 0 to prevent Decrementer exception after TSR is cleared
2. Write TBL to 0 to prevent Fixed Interval Timer and Watchdog Timer exceptions after TSR is cleared, and to prevent increment into TBH prior to full initialization
3. CCR1[TCS] (Timer Clock Select) can be initialized here, or earlier with the rest of the CCR1.
4. Clear TSR to clear all timer exception status
5. Write TCR to configure and enable timers as desired

Software must take care with respect to the enabling of the Watchdog Timer reset function, as once this function is enabled, it cannot be disabled except by reset itself

6. Initialize TBH value as desired
7. Initialize TBL value as desired
8. Initialize DECAR to desired value (if enabling the auto-reload function)
9. Initialize DEC to desired value

#### 14. Initialize facilities outside the processor core which are possible sources of asynchronous interrupt requests (including DCRs and/or other memory-mapped resources)

This must be done prior to enabling asynchronous interrupts in the MSR

15. Initialize the MSR to enable interrupts as desired
  1. Set MSR[CE] to enable/disable Critical Input and Watchdog Timer interrupts
  2. Set MSR[EE] to enable/disable External Input, Decrementer, and Fixed Interval Timer interrupts
  3. Set MSR[DE] to enable/disable Debug interrupts
  4. Set MSR[ME] to enable/disable Machine Check interrupts

Software should first check the status of the ESR[MCI] field and MCSR[MCS] field to determine whether any Machine Check exceptions have occurred after these fields were cleared by reset and before Machine Check interrupts were enabled (by this step). Any such exceptions would have set ESR[MCI] or MCSR[MCS] to 1, and this status can only be cleared explicitly by software. After the MCSR[MCS] field is known to be clear, the MCSR status bits (MCSR[1:8]) should be cleared by software to avoid possible confusion upon later service of a machine check interrupt. Once MSR[ME] has been set to 1, subsequent Machine Check exceptions will result in a Machine Check interrupt.

5. Context synchronize to establish new MSR context (**isync**)
16. Initialize any other processor core resources as required by the system (GPRs, SPRGs, and so on)
17. Initialize any other facilities outside the processor core as required by the system
18. Initialize system memory as required by the system software

Synchronize any program memory changes as required. (See *Self-Modifying Code* on page 103 for more information on the instruction sequence necessary to synchronize changes to program memory prior to executing the new instructions)

19. Start the system software

System software is generally responsible for initializing and/or managing the rest of the MSR fields, including:

1. MSR[FP] to enable or disable the execution of floating-point instructions
2. MSR[FE0,FE1] to enable/disable Floating-Point Enabled exception type Program interrupts
3. MSR[PR] to specify user mode or supervisor mode
4. MSR[IS,DS] to specify application address space or system address space for instructions and data
5. MSR[WE] to place the processor into Wait State (halt execution pending an interrupt)



## 4. Instruction and Data Caches

The PPC440x5 core provides separate instruction and data cache controllers and arrays, which allow concurrent access and minimize pipeline stalls. The storage capacity of the cache arrays, which can range from 8KB–32KB each, depends upon the implementation. Both cache controllers have 32-byte lines, and both are highly associative, having 64-way set-associativity for 32KB and 16KB sizes, and 32-way set-associativity for the 8KB size. The PowerPC instruction set provides a rich set of cache management instructions for software-enforced coherency. The PPC440x5 implementation also provides special debug instructions that can directly read the tag and data arrays. The cache controllers interface to the processor local bus (PLB) for connection to the IBM CoreConnect system-on-a-chip environment.

Both the data and instruction caches are parity protected against soft errors. If such errors are detected, the CPU will vector to the machine check interrupt handler, where software can take appropriate action. The details of suggested interrupt handling are described below in *Section 4.2 Instruction Cache Controller* and in *Section 4.3 Data Cache Controller*

The rest of this chapter provides more detailed information about the operation of the instruction and data cache controllers and arrays.

### 4.1 Cache Array Organization and Operation

The instruction and data cache arrays are organized identically, although the fields of the tag and data portions of the arrays are slightly different because the functions of the arrays differ, and because the instruction cache is virtually tagged while the data cache has real tags.

The associativity of each cache varies according to its size: the 32KB and 16KB cache sizes are 64-way set-associative, while the 8KB cache size is 32-way set-associative. Accordingly, the number of “sets” in each cache varies according to its size: the 32KB cache has 16 sets, while the 16KB and 8KB caches have 8 sets. Regardless of cache array size, the cache line size is always 32 bytes.

The organization of the cache into “ways” and “sets” is as follows. Using the 32KB cache as an example, there are 64 ways in each set, with a set consisting of all 64 lines (one line from each way) at which a given memory location can reside. Conversely, and again using the 32KB cache as an example, there are 16 sets in each way, with a way consisting of 16 lines (one from each set).

Table 4-1 on page 94 illustrates generically the ways and sets of the cache arrays, for any cache size, while Table 4-2 on page 94 provides specific values for the parameters used in *Table 4-1*, for the different cache sizes. As shown in *Table 4-2*, the tag field for each line in each way holds the high-order address bits associ-

ated with the line that currently resides in that way. The middle-order address bits form an index to select a specific set of the cache, while the five lowest-order address bits form a byte-offset to choose a specific byte (or bytes, depending on the size of the operation) from the 32-byte cache line.

Table 4-1. Instruction and Data Cache Array Organization

	Way 0	Way 1	• • •	Way $w - 2$	Way $w - 1$
Set 0	Line 0	Line $n$	• • •	Line $(w - 2)n$	Line $(w - 1)n$
Set 1	Line 1	Line $n + 1$	• • •	Line $(w - 2)n + 1$	Line $(w - 1)n + 1$
•	•	•	•	•	•
•	•	•	•	•	•
•	•	•	•	•	•
Set $n - 2$	Line $n - 2$	Line $2n - 2$	• • •	Line $(w - 1)n - 2$	Line $wn - 2$
Set $n - 1$	Line $n - 1$	Line $2n - 1$	• • •	Line $(w - 1)n - 1$	Line $wn - 1$

Table 4-2. Cache Sizes and Parameters

Array Size	$w$ (Ways)	$n$ (Sets)	Tag Address Bits <sup>1</sup>	Set Address Bits	Byte Offset Address Bits
8KB	32	8	$A_{0:23}$	$A_{24:26}$	$A_{27:31}$
16KB	64	8	$A_{0:23}$	$A_{24:26}$	$A_{27:31}$
32KB	64	16	$A_{0:22}$	$A_{23:26}$	$A_{27:31}$

1. The tag address bits shown in the table refer to the effective address bits, and are for illustrative purposes only. Because the instruction cache is tagged with the virtual address, and the data cache is tagged with the real address, the actual tag address bits contained within each array are different. See Figure 4-8 and Figure 4-9 on page 110 for instruction cache tag information, and Figure 4-10 and Figure 4-11 on page 124 for data cache tag information. Also, see Section 4.2.3.2 on page 104 for details on instruction cache synonyms associated with the use of virtual tags for the instruction cache.

### 4.1.1 Cache Line Replacement Policy

Memory addresses are specified as being cacheable or caching inhibited on a page basis, using the caching inhibited (I) storage attribute (see *Caching Inhibited (I)* on page 141). When a program references a cacheable memory location and that location is not already in the cache (a *cache miss*), the line may be brought into the cache (a *cache line fill* operation) and placed into any one of the ways within the set selected by the middle portion of the address (the specific address bits that select the set are specified in Table 4-2). If the particular way within the set already contains a valid line from some other address, the existing line is removed and replaced by the newly referenced line from memory. The line being replaced is referred to as the *victim*.

The way selected to be the victim for replacement is controlled by a field within a Special Purpose Register (SPR). There is a separate "victim index field" for each set within the cache. The registers controlling the victim selection are shown in Figure 4-1.




Figure 4-1. Instruction Cache Normal Victim Registers (INV0–INV3) Instruction Cache Transient Victim Registers (ITV0–ITV3) Data Cache Normal Victim Registers (DNV0–DNV3) Data Cache Transient Victim Registers (DTV0–DTV3)

0:7	VNDXA	Victim Index A (for cache lines with EA[25:26] = 0b00)	For all victim index fields, the number of bits used to select the cache way for replacement depends on the implemented cache size. See <i>Table 4-5</i> on page 117. for more information.
8:15	VNDXB	Victim Index B (for cache lines with EA[25:26] = 0b01)	
16:23	VNDXC	Victim Index C (for cache lines with EA[25:26] = 0b10)	
24:31	VNDXD	Victim Index D (for cache lines with EA[25:26] = 0b11)	

Each of the 16 SPRs illustrated in *Figure 4-1* can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**. In general, however, these registers are initialized by software once at startup, and then are managed automatically by hardware after that. Specifically, every time a new cache line is placed into the cache, the appropriate victim index field (as controlled by the type of access and the particular cache set being updated) is first referenced to determine which way within that set should be replaced. Then, that same field is incremented such that the ways within that set are replaced in a *round-robin* fashion as each new line is brought into that set. When the victim index field value reaches the index of the last way (according to the size of the cache and the type of access being performed), the value is *wrapped back* to the index of the first way for that type of access. The first and last ways for the different types of accesses are controlled by fields in a pair of *victim limit* SPRs, one for each cache (see *Cache Locking and Transient Mechanism* on page 96 for more information).

The size of the victim index fields varies according to the size of the respective cache. Also, which field is used varies according to the type of access, the size of the cache, and the address of the cache line. Table 4-3 describes the correlation between the victim index fields and different access types, cache sizes, and addresses.

Table 4-3. Victim Index Field Selection

Address <sub>23:26</sub> <sup>3</sup>	Victim Index Field <sup>1,2</sup>		
	8KB Cache	16KB Cache	32KB Cache
0	xxV0[VNDXA] <sub>3:7</sub>	xxV0[VNDXA] <sub>2:7</sub>	xxV0[VNDXA] <sub>2:7</sub>
1	xxV0[VNDXB] <sub>3:7</sub>	xxV0[VNDXB] <sub>2:7</sub>	xxV0[VNDXB] <sub>2:7</sub>
2	xxV0[VNDXC] <sub>3:7</sub>	xxV0[VNDXC] <sub>2:7</sub>	xxV0[VNDXC] <sub>2:7</sub>
3	xxV0[VNDXD] <sub>3:7</sub>	xxV0[VNDXD] <sub>2:7</sub>	xxV0[VNDXD] <sub>2:7</sub>
4	xxV1[VNDXA] <sub>3:7</sub>	xxV1[VNDXA] <sub>2:7</sub>	xxV1[VNDXA] <sub>2:7</sub>
5	xxV1[VNDXB] <sub>3:7</sub>	xxV1[VNDXB] <sub>2:7</sub>	xxV1[VNDXB] <sub>2:7</sub>
6	xxV1[VNDXC] <sub>3:7</sub>	xxV1[VNDXC] <sub>2:7</sub>	xxV1[VNDXC] <sub>2:7</sub>
7	xxV1[VNDXD] <sub>3:7</sub>	xxV1[VNDXD] <sub>2:7</sub>	xxV1[VNDXD] <sub>2:7</sub>
8	xxV0[VNDXA] <sub>3:7</sub>	xxV0[VNDXA] <sub>2:7</sub>	xxV2[VNDXA] <sub>2:7</sub>
9	xxV0[VNDXB] <sub>3:7</sub>	xxV0[VNDXB] <sub>2:7</sub>	xxV2[VNDXB] <sub>2:7</sub>
10	xxV0[VNDXC] <sub>3:7</sub>	xxV0[VNDXC] <sub>2:7</sub>	xxV2[VNDXC] <sub>2:7</sub>
11	xxV0[VNDXD] <sub>3:7</sub>	xxV0[VNDXD] <sub>2:7</sub>	xxV2[VNDXD] <sub>2:7</sub>
12	xxV1[VNDXA] <sub>3:7</sub>	xxV1[VNDXA] <sub>2:7</sub>	xxV3[VNDXA] <sub>2:7</sub>
13	xxV1[VNDXB] <sub>3:7</sub>	xxV1[VNDXB] <sub>2:7</sub>	xxV3[VNDXB] <sub>2:7</sub>
14	xxV1[VNDXC] <sub>3:7</sub>	xxV1[VNDXC] <sub>2:7</sub>	xxV3[VNDXC] <sub>2:7</sub>
15	xxV1[VNDXD] <sub>3:7</sub>	xxV1[VNDXD] <sub>2:7</sub>	xxV3[VNDXD] <sub>2:7</sub>

**Note 1:** In the victim index field columns, the “xx” in the SPR name refers to one of “IN”, “IT”, “DN”, or “DT”, depending on whether the access is to the instruction or data cache, and whether it is a “normal” or a “transient” access (See *Cache Locking and Transient Mechanism* on page 96.)

**Note 2:** As shown in the table, the 8KB cache size only uses bits 3:7 of the victim index fields to select a way, since there are only 32 ways. Similarly, the 16KB and 32KB cache sizes use bits 2:7 of the victim index fields, since those cache sizes have 64 ways. In all cases, the unused bits of the victim index fields are reserved. The size of the fields of the victim limit registers (IVLIM, DVLIM) are similarly affected by the number of sets in the cache (See *Cache Locking and Transient Mechanism* on page 96.)

**Note 3:** Since the 8KB and 16KB cache sizes only have 8 sets, they only use Address<sub>24:26</sub> to select the set and the victim index field, and thus they do not use the xxV2 and xxV3 SPRs.

#### 4.1.2 Cache Locking and Transient Mechanism


Both caches support locking, at a “way” granularity. Any number of ways can be locked, from 0 ways to one less than the total number of ways (64 ways for 32KB and 16KB cache sizes, 32 ways for the 8KB cache size). At least one way must always be left unlocked, for use by cacheable line fills. Each way contains one line from each set; that is, either 16 lines (512 bytes), for the 32KB cache size, or 8 lines (256 bytes), for the 16KB and 8KB cache sizes.

In addition, a portion of each cache can be designated as a “transient” region, by specifying that only a limited number of ways are used for cache lines from memory pages that are identified as being transient in nature by a storage attribute from the MMU (see *Memory Management* on page 129). For the instruction cache,



such memory pages can be used for code sequences that are unlikely to be reused once the processor moves on to the next series of instruction lines. Thus, performance may be improved by preventing each series of instruction lines from overwriting the rest of the “regular” code in the instruction cache. Similarly, for the data cache, transient pages can be used for large “streaming” data structures, such as multimedia data. As each piece of the data stream is processed and written back to memory, the next piece can be brought in, overwriting the previous (now obsolete) cache lines instead of displacing other areas of the cache, which may contain other data that should remain in the cache.

A set of fields in a pair of victim limit registers specifies which ways of the cache are used for normal accesses and/or transient accesses, as well as which ways are locked. These registers, Instruction Cache Victim Limit (IVLIM) and Data Cache Victim Limit (DVLIM), are illustrated in *Figure 4-2*. They can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.



*Figure 4-2. Instruction Cache Victim Limit (IVLIM) Data Cache Victim Limit (DVLIM)*

0:1		Reserved	
2:9	TFLOOR	Transient Floor	The number of bits in the TFLOOR field varies, depending on the implemented cache size. See <i>Table 4-5</i> on page 117 for more information.
10:12		Reserved	
13:20	TCEILING	Transient Ceiling	The number of bits in the TCEILING field varies, depending on the implemented cache size. See <i>Table 4-5</i> on page 117 for more information.
21:23		Reserved	
24:31	NFLOOR	Normal Floor	The number of bits in the NFLOOR field varies, depending on the implemented cache size. See <i>Table 4-5</i> on page 117 for more information.

When a cache line fill occurs as the result of a normal memory access (that is, one *not* marked as transient using the U1 storage attribute from the MMU; see *Memory Management* on page 129), the cache line to be replaced is selected by the corresponding victim index field from one of the normal victim index registers (INV0–INV3 for instruction cache lines, DNV0–DNV3 for data cache lines). As the processor increments any of these normal victim index fields according to the round-robin mechanism described in *Cache Line Replacement Policy* on page 94, the values of the fields are constrained to lie within the range specified by the NFLOOR field of the corresponding victim limit register, and the last way of the cache (way 31 for the 8KB cache size, way 63 for the 16KB or 32KB cache size). That is, when one of the normal victim index fields is incremented past the last way of the cache, it wraps back to the value of the NFLOOR field of the associated victim limit register.

Similarly, when a cache line fill occurs as the result of a transient memory access, the cache line to be replaced is selected by the corresponding victim index field from one of the transient victim index registers (ITV0–ITV3 for instruction cache lines, DTV0–DTV3 for data cache lines). As the processor increments any of these transient victim index fields according to the round-robin replacement mechanism, the values of the fields are constrained to lie within the range specified by the TFLOOR and the TCEILING fields of the corre-

sponding victim limit register. That is, when one of the transient victim index fields is incremented past the TCEILING value of the associated victim limit register, it wraps back to the value of the TFLOOR field of that victim limit register.

Given the operation of this mechanism, if both the NFLOOR and TFLOOR fields are set to 0, and the TCEILING is set to the index of the last way of the cache, then all cache line fills—both normal and transient—are permitted to use the entire cache, and nothing is locked. Alternatively, if both the NFLOOR and TFLOOR fields are set to values greater than 0, the lines in those ways of the cache whose indexes are between 0 and the lower of the two floor values are effectively *locked*, as no cache line fills (neither normal nor transient) will be allowed to replace the lines in those ways. Yet another example is when the TFLOOR is lower than the NFLOOR, and the TCEILING is lower than the last way of the cache. In this scenario, the ways between the TFLOOR and the NFLOOR contain only transient lines, while the ways between the NFLOOR and the TCEILING may contain either normal or transient lines, and the ways from the TCEILING to the last way of the cache contain only normal lines.

**Programming Note:** It is a programming error for software to program the TCEILING field to a value lower than that of the TFLOOR field. Furthermore, software must initialize each of the normal and transient victim index fields to values that are between the ranges designated by the respective victim limit fields, prior to performing any cacheable accesses intended to utilize these ranges.

In order to setup a locked area within the data cache, software must perform the following steps (the procedure for the instruction cache is similar, with **icbt** instructions substituting for **dcbt** instructions):

1. Execute **msync** and then **isync** to guarantee all previous cache operation have completed.
2. Mark all TLB entries associated with memory pages which are being used to perform the locking function as caching-inhibited. Leave the TLB entries associated with the memory pages containing the data which is to be locked into the data cache marked as cacheable, however.
3. Execute **msync** and then **isync** again, to cause the new TLB entry values to take effect.
4. Set both the NFLOOR and the TFLOOR values to the index of the first way which should be locked, and set the TCEILING value to the last way of the cache.
5. Set each of the normal and transient victim index fields to the same value as the NFLOOR and TFLOOR.
6. Execute **dcbt** instructions to the cache lines within the cacheable memory pages which contain the data which is to be locked in the data cache. The number of **dcbt** instructions executed to any given set should not exceed the number of ways which will exist in the locked region (otherwise not all of the lines will be able to be simultaneously locked in the data cache). Remember that when a series of **dcbt** instructions are executed to sequentially increasing addresses (with the address increment being the size of a cache block -- 32 bytes), it takes sixteen such **dcbt** operations (one for each set) before the next way of the initial set will be targeted again.
7. Execute **msync** and then **isync** again, to guarantee that all of the **dcbt** operations have completed and updated the corresponding victim index fields.
8. Set the NFLOOR, TFLOOR, and TCEILING values to the desired indices for the operating normal and transient regions of the cache. Both the NFLOOR and the TFLOOR values should be set higher than the highest locked way of the data cache; otherwise, subsequent normal and/or transient accesses could overwrite a way containing a line which was to be locked.
9. Set each of the normal and transient victim index fields to the value of the NFLOOR and TFLOOR, respectively.

10. Restore the cacheability of the memory pages which were used to perform the locking function to the desired operating values, by clearing the caching-inhibited attribute of the TLB entries which were updated in step 2.
11. Execute **msync** and then **isync** again, to cause the new TLB entry values to take effect.

The ways of the data cache whose indices are below the lower of the NFLOOR and TFLOOR values will now be locked.

Figure 4-3 and Figure 4-4 illustrate two of these examples of the use of the locking and transient mechanisms. Other configurations are possible, given the ability to program each of the victim limit fields to different relative values, although some configurations are not necessarily useful or practical.

Figure 4-3. Cache Locking and Transient Mechanism (Example 1)<sup>1</sup>

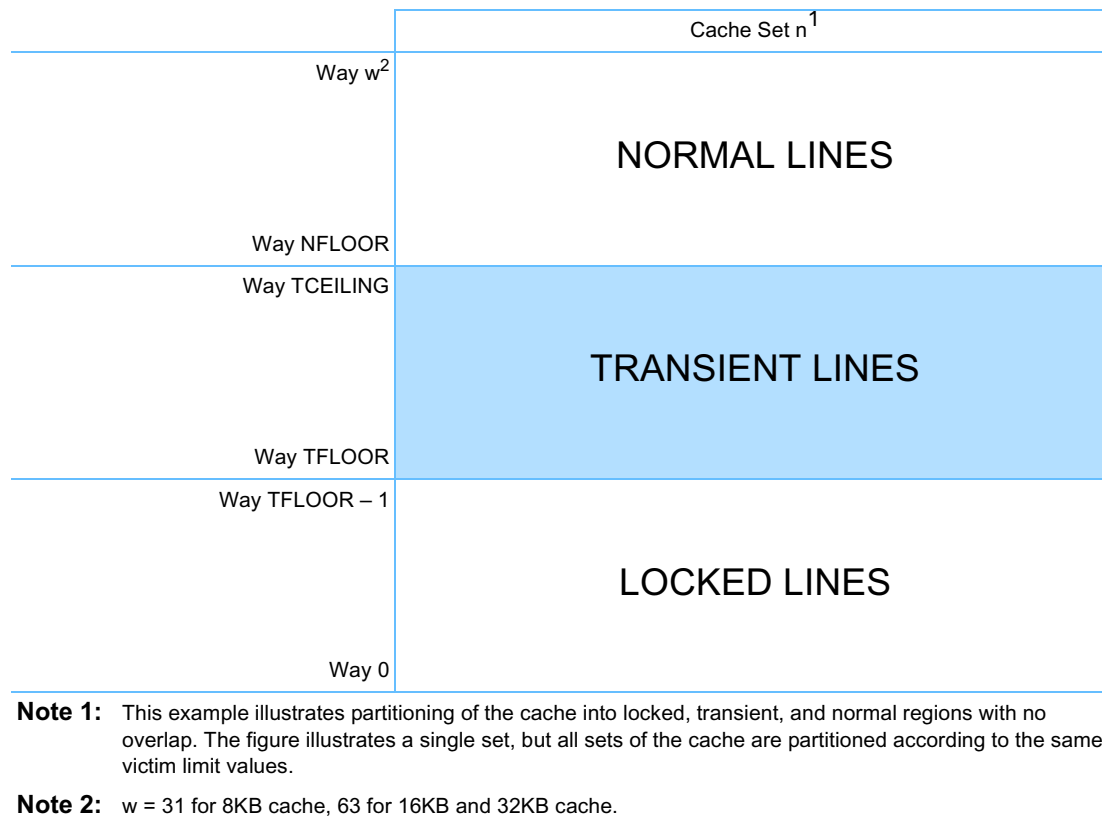
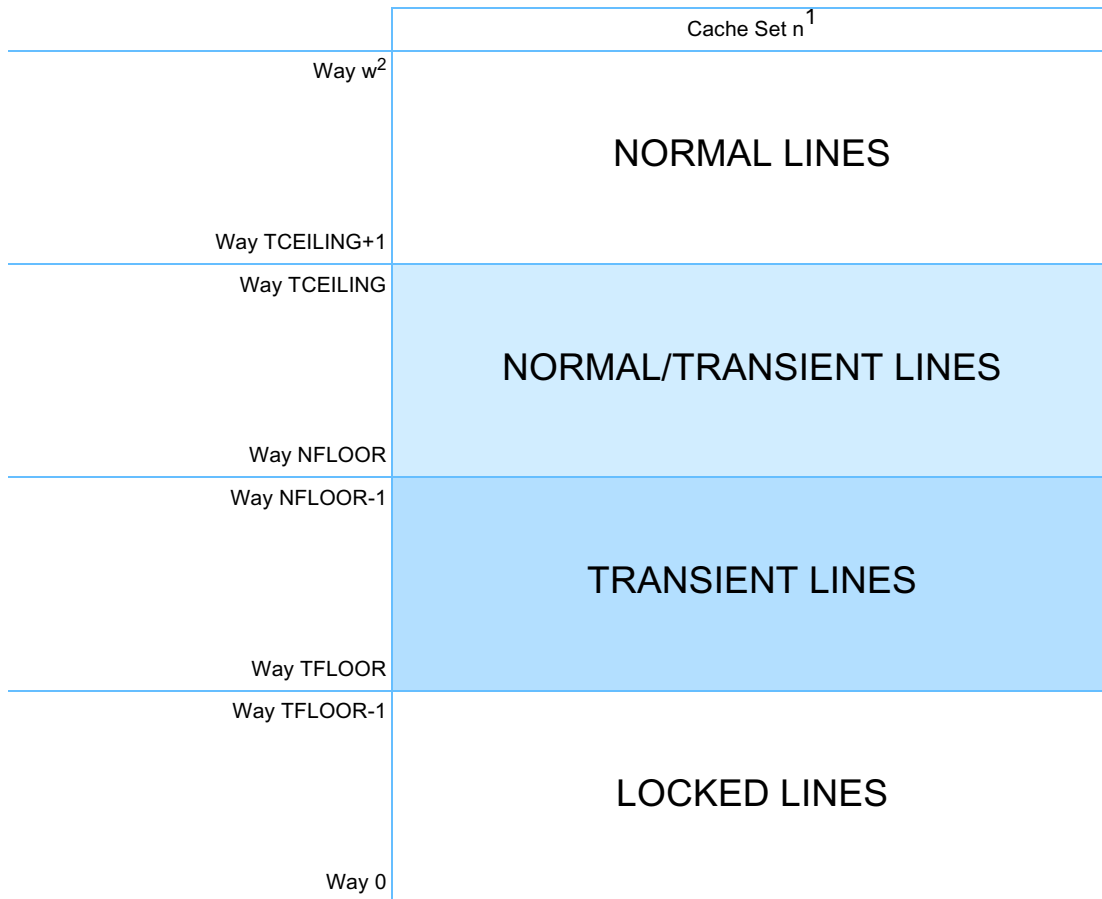


Figure 4-4. Cache Locking and Transient Mechanism (Example 2)



**Note 1:** This example illustrates partitioning of the cache into locked, transient, and normal regions where the transient and normal regions partially overlap. The figure illustrates a single set, but all sets of the cache are partitioned according to the same victim limit values.

**Note 2:**  $w = 31$  for 8KB cache, 63 for 16KB and 32KB cache.

## 4.2 Instruction Cache Controller

The instruction cache controller (ICC) delivers two instructions per cycle to the instruction unit of the PPC440x5 core. The ICC interfaces to the PLB using a 128-bit read interface, although it supports direct attachment to 32-bit and 64-bit PLB subsystems, as well as 128-bit PLB subsystems. The ICC handles frequency synchronization between the PPC440x5 core and the PLB, and can operate at any ratio of  $n:1$ ,  $n:2$ , and  $n:3$ , where  $n$  is an integer greater than the corresponding denominator.

The ICC provides a speculative prefetch mechanism which can be configured to automatically prefetch a burst of up to three additional lines upon any fetch request which misses in the instruction cache.

The ICC also handles the execution of the PowerPC instruction cache management instructions, for touching (prefetching) or invalidating cache lines, or for flash invalidation of the entire cache. Resources for controlling and debugging the instruction cache operation are also provided.

The rest of this section describes each of these functions in more detail.

#### 4.2.1 ICC Operations

When the ICC receives an instruction fetch request from the instruction unit of the PPC440x5 core, the ICC simultaneously searches the instruction cache array for the cache line associated with the virtual address of the fetch request, and translates the virtual address into a real address (see *Memory Management* on page 129 for information about address translation). If the requested cache line is found in the array (a cache *hit*), the pair of instructions at the requested address are returned to the instruction unit. If the requested cache line is *not* found in the array (a cache *miss*), the ICC sends a request for the entire cache line (32 bytes) to the instruction PLB interface, using the real address. Note that the entire 32-byte cache line is requested, even if the caching inhibited (I) storage attribute is set for the memory page containing that cache line (see *Caching Inhibited (I)* on page 141). Also note that the request to the instruction PLB interface is sent using the specific instruction address requested by the instruction unit, so that the memory subsystem may read the cache line *target word first* (if it supports such operation) and supply the requested instructions before retrieving the rest of the cache line.

As the ICC receives each portion of the cache line from the instruction PLB interface, it is placed into the instruction cache line fill data (ICLFD) buffer. Instructions from this buffer may be *bypassed* to the instruction unit as requested, without waiting for the entire cache line to be filled. Once the entire cache line has been filled into the buffer, and assuming that the memory page containing that line is cacheable, it is written into the instruction cache. If the memory page containing the line is caching inhibited, the line will remain in the ICLFD until it is displaced by a subsequent request for another cache line (either cacheable or caching inhibited).

If a memory subsystem error (such as an address time-out, invalid address, or some other type of hardware error external to the PPC440x5 core) occurs during the filling of the cache line, the line will not be written into the instruction cache, although instructions from the line may still be forwarded to the instruction unit from the ICLFD. Later, if execution of an instruction from that line is attempted, an Instruction Machine Check exception will be reported, and a Machine Check interrupt (if enabled) will result. See *Machine Check Interrupt* on page 172 for more information on Machine Check interrupts.

Once a request for a cache line read has been requested on the instruction PLB interface, the entire line read will be performed and the line will be written into the instruction cache (assuming no error occurs on the read), regardless of whether or not the instruction stream branches (or is interrupted) away from the line being read. This behavior is due to the nature of the PLB architecture, and the fact that once started, a cache line read request type cannot be abandoned. This does not mean, however, that the ICC will wait for this cache line read to complete before responding to a new request from the instruction unit (due, perhaps, to a branch redirection, or an interrupt). Instead, the ICC will immediately access the cache to determine if the cache line at the new address requested by the instruction unit is already in the cache. If so, the requested pair of instructions from this line will immediately be forwarded to the instruction unit, while the ICC in parallel continues to fill the previously requested cache line. In other words, the instruction cache is completely *non-blocking*.

If the newly requested cache line is instead a miss in the instruction cache, the ICC will immediately attempt to cancel the previous cache line read request. If the previous cache line read request has not yet been requested on the PLB bus, the old request will be cancelled and the new request will be made. If the previous cache line read request has already been requested, then as previously stated it cannot be abandoned, but the ICC will immediately present the request for the new cache line, such that it may be serviced immediately after the previous cache line read is completed. The ICC never aborts any PLB request once it has been made, except when a processor reset occurs while the PLB request is being made.

**Programming Note:**

It is a programming error for an instruction fetch request to reference a valid cache line in the instruction cache if the caching inhibited storage attribute is set for the memory page containing the cache line. The result of attempting to execute an instruction from such an access is undefined. After processor reset, hardware automatically sets the caching inhibited storage attribute for the memory page containing the reset address, and also automatically flush invalidates the instruction cache. Subsequently, lines will not be placed into the instruction cache unless they are accessed by reference to a memory page for which the caching inhibited attribute has been turned off. If software subsequently turns on the caching inhibited storage attribute for such a page, software must make sure that no lines from that page remain valid in the instruction cache, before attempting to fetch and execute instructions from the (now caching inhibited) page.

**4.2.2 Speculative Prefetch Mechanism**

The ICC can be configured to automatically prefetch up to three more cache lines upon (in addition to the line being requested by the instruction unit) in response to a cache miss. This speculative prefetch only occurs on requests for lines from cacheable memory pages, and then only if enabled by the setting of certain fields in the Core Configuration Register 0 (CCR0) (see *Figure 4-5* on page 106).

CCR0[ICSLC] specifies the number of additional cache lines (from 0 to 3) to speculatively prefetch upon an instruction cache miss. If this field is non-zero, upon an instruction cache miss, the ICC will first check the cache to see whether the additional lines are themselves already in the cache. If not, then the ICC will present a fixed-length burst request to the instruction PLB interface, requesting the additional cache line(s). The burst request is presented after the cache line request for the initial cache line requested by the instruction unit is presented and acknowledged on the PLB.

The speculative line fill mechanism will not request lines past the end of the minimum memory page size, which is 1KB. That is, if the line requested by the instruction unit is at or near the end of an aligned 1KB boundary, the speculative prefetch mechanism will only request those additional lines specified by the CCR0[ICSLC] field that are also within the same 1KB page of memory. This allows the speculative prefetch mechanism to operate without having to access the Memory Management Unit (MMU) for a translation for the next page address.

Another field in the CCR0 register, CCR0[ICSLT], specifies a *threshold* value that is used to determine whether the speculative burst request should be abandoned prior to completion, as a result of a change in direction in the instruction stream (such as a branch or interrupt). If the instruction unit requests a new cache line and the new request is a hit in the instruction cache, both the original line fill request and any speculative burst request associated with it will be unaffected. Furthermore, if the new cache line requested by the instruction unit is a miss in the instruction cache, any prior request which has not yet been requested on the PLB interface will be cancelled, regardless of the value of CCR0[ICSLT]. However, if a prior speculative burst request has already been requested on the PLB interface, the value of CCR0[ICSLT] determines if and when the speculative burst request will be abandoned. CCR0[ICSLT] specifies the number of *doublewords* (8-byte units) of the *current* cache line which must already have been received by the ICC, in order that the filling of the current cache line will *not* be abandoned (note that in this context, the term “current” refers to the cache line with which the next PLB data transfer is associated, at the time that the ICC determines that it needs to request a new line). That is, if the ICC has already received the number of doublewords indicated by CCR0[ICSLT], the ICC will not terminate the burst until it has received that entire cache line. All additional lines beyond the one in progress at the time that the ICC determines that it needs to request a new line *will* be abandoned. For example, if CCR0[ICSLC] is set to 3, and the ICC is in the middle of receiving the data for the

first of the three speculative lines at the time that the new instruction cache miss request is received from the instruction unit, the second and third lines of the speculative burst will be abandoned, and whether the first of the speculative lines is abandoned is controlled by CCR0[ICSLT].

Since cache lines contain 32 bytes, there are four doublewords in each cache line. Thus, CCR0[ICSLT] can be set to a value from 0 to 3. If CCR0[ICSLT] = 0, the current line fill will be completed regardless of how many doublewords have already been received. Similarly, if CCR0[ICSLT] = 3, the current line fill will be abandoned if only two or fewer doublewords have been received by the ICC.

If at the time that the ICC determines that it needs to request a new line and abandon a speculative burst request, the ICC has still not received all of the data associated with the initial cache line request which prompted the speculative burst request, then this initial cache line is considered the “current” line, and the speculative burst request will be abandoned without filling *any* of the speculative lines, regardless of the setting of CCR0[ICSLT]. The filling of the initial cache line *will* be completed, however, as the PLB protocol does not provide for the abandonment of the cache line (non-burst) request type.

Regardless of the value of CCR0[ICSLT], any time that a cache line fill is abandoned such that all of the data for that cache line is not received, the line may still be used to bypass instructions to the instruction unit, but it will not be written into the instruction cache, and it will be overwritten in the ICLFD buffer as soon as instructions for a new line begin arriving from the PLB.

### 4.2.3 Instruction Cache Coherency

In general, the PPC440x5 does not automatically enforce coherency between the instruction cache, data cache, and memory. If the contents of memory location are changed, either within the data cache or within memory itself, and whether by the PPC440x5 core through the execution of store instructions or by some other mechanism in the system writing to memory, software must use cache management instructions to ensure that the instruction cache is made coherent with these changes. This involves invalidating any obsolete copies of these memory locations within the instruction cache, so that they will be reread from memory the next time they are referenced by program execution.

#### 4.2.3.1 Self-Modifying Code

To illustrate the use of the cache management instructions to enforce instruction cache coherency, consider the example of *self-modifying code*, whereby the program executing on the PPC440x5 core stores new data to memory, with the intention of later branching to and executing this new “data,” which are actually instructions.

The following code example illustrates the required sequence for software to use when writing self-modifying code. This example assumes that *addr1* references a cacheable memory page.

```

stw      regN, addr1    # store the data (an instruction) in regN to addr1 in the data cache
dcbst   addr1          # write the new instruction from the data cache to memory
msync                   # wait until the data actually reaches the memory
icbi    addr1          # invalidate addr1 in the instruction cache if it exists
msync   addr1          # wait for the instruction cache invalidation to take effect
isync                   # flush any prefetched instructions within the ICC and instruction
                        # unit and re-fetch them (an older copy of the instruction at addr1
                        # may have already been fetched)
  
```

At this point, software may begin executing the instruction at *addr1* and be guaranteed that the new instruction will be recognized.



### 4.2.3.2 Instruction Cache Synonyms

A synonym is a cache line that is associated with the same real address as another cache line that is in the cache array at the same time. Such synonyms can occur when different virtual addresses are mapped to the same real address, and the virtual address is used either as an index to the cache array (a *virtually-indexed* cache) or as the cache line tag (a *virtually-tagged* cache).

The instruction cache on the PPC440x5 is real-indexed but virtually-tagged and thus it is possible for synonyms to exist in the cache. (The data cache on the other hand is both real-indexed and real-tagged, and thus cannot have any synonyms.) Because of this, special care must be taken when managing instruction cache coherency and attempting to invalidate lines in the cache.

As explained in *Memory Management* on page 129, the virtual address (VA) consists of the 32-bit effective address (EA; for instruction fetches, this is the address calculated by the instruction unit and sent to the ICC) combined with the 8-bit Process ID (PID) and the 1-bit address space (MSR[IS] for instruction fetches). As described in *Table 4-2* on page 94, and using the 32KB cache as an example,  $VA_{27:31}$  chooses the byte offset within the cache line, while  $VA_{23:26}$  is used as the *index* to select a set, and then the rest of the virtual address is used as the *tag*. The tag thus consists of  $EA_{0:22}$ , the PID, and MSR[IS] (for instruction fetches; for cache management instructions such as **icbi**, MSR[DS] is used to specify the address space; see the instruction descriptions for the instruction cache management instructions for more information). The tag portion of the VA is compared against the corresponding tag fields of each cache line within the way selected by  $VA_{23:26}$ .

Note that the address translation architecture of PowerPC Book-E is such that the low-order address bits 22:31 are always the same for the EA, VA, and real address (RA), because these bits are never translated due to the minimum page size being 1KB (these low-order 10 bits are always used for the byte offset within the page). As the page size increases, more and more low-order bits are used for the byte offset within the page, and thus fewer and fewer bits are translated between the VA and the RA (see *Table 5-3* on page 137). Synonyms only become possible when the system-level memory management software establishes multiple mappings to the same real page, which by definition involves different virtual addresses (either through differences in the higher-order EA bits which make up the VA, or through different process IDs, or different address spaces, or some combination of these three portions of the VA).

A further requirement for synonyms to exist in the instruction cache is for more than one of the virtual pages which map to a given real page to have *execute permission*, and for these pages to be cacheable (cache lines associated with pages without execute permission, or for which the caching inhibited storage attribute is set, cannot be placed in the instruction cache).

If the system-level memory management software permits instruction cache synonyms to be created, then extra care must be taken when attempting to invalidate instruction cache lines associated with a particular address. If software desires to invalidate only the cache line which is associated with a specific VA, then only a single **icbi** instruction need be executed, specifying that VA. If, however, software wishes to invalidate *all* instruction cache lines which are associated with a particular RA, then software must issue an **icbi** instruction for *each* VA which has a mapping to that particular RA and for which a line might exist in the instruction cache. In order to do this, the memory management software must keep track of which mappings to a given RA exist (or ever existed, if a mapping has been removed but cache lines associated with it might still exist), so that **icbi** instructions can be executed using the necessary VAs.

Alternatively, software can execute an **iccci** instruction, which flash invalidates the entire instruction cache without regard to the addresses with which the cache lines are associated.



#### 4.2.4 Instruction Cache Control and Debug

The PPC440x5 core provides various registers and instructions to control instruction cache operation and to help debug instruction cache problems.

##### 4.2.4.1 Instruction Cache Management and Debug Instruction Summary

For detailed descriptions of the instructions summarized in this section, see *Instruction Set* on page 243. Also, see *Instruction Cache Coherency* on page 103 for more information on how these instructions are used to manage coherency in the instruction cache.

In the instruction descriptions, the term “block” describes the unit of storage operated on by the cache block instructions. For the PPC440x5 core, this is the same as a cache line.

The following instructions are used by software to manage the instruction cache:

<b>icbi</b>	Instruction Cache Block Invalidate  Invalidates a cache block.
<b>icbt</b>	Instruction Cache Block Touch  Initiates a block fill, enabling a program to begin a cache block fetch before the program needs an instruction in the block. The program can subsequently branch to the instruction address and fetch the instruction without incurring a cache miss.  See <i>icbt Operation</i> on page 108.
<b>iccci</b>	Instruction Cache Congruence Class Invalidate  Flash invalidates the entire instruction cache. Execution of this instruction is privileged.
<b>icread</b>	Instruction Cache Read  Reads a cache line (tag and data) from a specified index of the instruction cache, into a set of SPRs. Execution of this instruction is privileged.  See <i>icread Operation</i> on page 109.

##### 4.2.4.2 Core Configuration Register 0 (CCR0)

The CCR0 register controls the speculative prefetch mechanism and the behavior of the **icbt** instruction. The CCR0 register also controls various other functions within the PPC440x5 core that are unrelated to the instruction cache. Each of these functions is discussed in more detail in the related sections of this manual.

Figure 4-5 illustrates the fields of the CCR0 register.




Figure 4-5. Core Configuration Register 0 (CCR0)

0		Reserved	
1	PRE	Parity Recoverability Enable 0 Semi-recoverable parity mode enabled for data cache 1 Fully recoverable parity mode enabled for data cache	Must be set to 1 to guarantee full recoverability from MMU and data cache parity errors.
2:3		Reserved	
4	CRPE	Cache Read Parity Enable 0 Disable parity information reads 1 Enable parity information reads	When enabled, execution of <b>icread</b> , <b>dcread</b> , or <b>tlbre</b> loads parity information into the ICDBTRH, DCDBTRL, or target GPR, respectively.
5:9		Reserved	
10	DSTG	Disable Store Gathering 0 Enabled; stores to contiguous addresses may be gathered into a single transfer 1 Disabled; all stores to memory will be performed independently	See <i>Store Gathering</i> on page 116.
11	DAPUIB	Disable APU Instruction Broadcast 0 Enabled. 1 Disabled; instructions not broadcast to APU for decoding	This mechanism is provided as a means of reducing power consumption when an auxilliary processor is not attached and/or is not being used. See <i>Initialization</i> on page 83.
12:15		Reserved	
16	DTB	Disable Trace Broadcast 0 Enabled. 1 Disabled; no trace information is broadcast.	This mechanism is provided as a means of reducing power consumption when instruction tracing is not needed. See <i>Initialization</i> on page 83.
17	GICBT	Guaranteed Instruction Cache Block Touch 0 <b>icbt</b> may be abandoned without having filled cache line if instruction pipeline stalls. 1 <b>icbt</b> is guaranteed to fill cache line even if instruction pipeline stalls.	See <i>icbt Operation</i> on page 108.
18	GDCBT	Guaranteed Data Cache Block Touch 0 <b>dcbt/dcbtst</b> may be abandoned without having filled cache line if load/store pipeline stalls. 1 <b>dcbt/dcbtst</b> are guaranteed to fill cache line even if load/store pipeline stalls.	See <i>Data Cache Control and Debug</i> on page 121.
19:22		Reserved	
23	FLSTA	Force Load/Store Alignment 0 No Alignment exception on integer storage access instructions, regardless of alignment 1 An alignment exception occurs on integer storage access instructions if data address is not on an operand boundary.	See <i>Load and Store Alignment</i> on page 114.
24:27		Reserved	

28:29	ICSLC	Instruction Cache Speculative Line Count	Number of additional lines (0–3) to fill on instruction fetch miss. See <i>Speculative Prefetch Mechanism</i> on page 102.
30:31	ICSLT	Instruction Cache Speculative Line Threshold	Number of doublewords that must have already been filled in order that the current speculative line fill is <i>not</i> abandoned on a redirection of the instruction stream. See <i>Speculative Prefetch Mechanism</i> on page 102.

#### 4.2.4.3 Core Configuration Register 1 (CCR1)

The CCR1 register controls parity error insertion for software testing, one option for line flush behavior in the D-cache, and a control bit that selects the timer input clock. Each of these functions is discussed in more detail in the related sections of this manual.

Figure 4-6 illustrates the fields of the CCR1 register.




Figure 4-6. Core Configuration Register 1 (CCR1)

0:7	ICDPEI	Instruction Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded when the instruction cache is filled. Each of the 8 bits corresponds to one of the instruction words in the line.
8:9	ICTPEI	Instruction Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the instruction cache.
10:11	DCTPEI	Data Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the data cache.
12	DCDPEI	Data Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the data field in the data cache.
13	DCUPEI	Data Cache U-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bit recorded for the U fields in the data cache.
14	DCMPEI	Data Cache Modified-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the modified (dirty) field in the data cache.
15	FCOM	Force Cache Operation Miss 0 normal operation 1 cache ops appear to miss the cache	Force <b>icbt</b> , <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>dcbi</b> , and <b>dcbz</b> to appear to miss the caches. The intended use is with <b>icbt</b> and <b>dcbt</b> only, which will fill a duplicate line and allow testing of multi-hit parity errors. See <i>Section 4.2.4.7 Simulating Instruction Cache Parity Errors for Software Testing</i> on page 111 and <i>Figure 4.3.3.7</i> on page 126.

16:19	MMUPEI	Memory Management Unit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the MMU.
20	FFF	Force Full-line Flush 0 flush only as much data as necessary. 1 always flush entire cache lines	When flushing 32-byte (8-word) lines from the data cache, normal operation is to write nothing, a double word, quad word, or the entire 8-word block to the memory as required by the dirty bits. This bit ensures that none or all dirty bits are set so that either nothing or the entire 8-word block is written to memory when flushing a line from the data cache. Refer to <i>Section 4.3.1.4 Line Flush Operations</i> on page 117.
21:23		Reserved	
24	TCS	Timer Clock Select 0 CPU timer advances by one at each rising edge of the CPU input clock (CPMC440CLOCK). 1 CPU timer advances by one for each rising edge of the CPU timer clock (CPMC440TIMERCLOCK).	When TCS = 1, CPU timer clock input can toggle at up to half of the CPU clock frequency.
25:31		Reserved	

#### 4.2.4.4 icbt Operation

The **icbt** instruction is typically used as a “hint” to the processor that a particular block of instructions is likely to be executed in the near future. Thus the processor can begin filling that block into the instruction cache, so that when the executing program eventually branches there the instructions will already be present in the cache, thereby improving performance.

Of course, it would not typically be advantageous if the filling of the cache line requested by the **icbt** itself caused a delay in the fetching of instructions needed by the currently executing program. For this reason, the default behavior of the **icbt** instruction is for it to have the lowest priority for sending a request to the PLB. If a subsequent instruction cache miss occurs due to a request from the instruction unit, then the line fill for the **icbt** will be abandoned (if it has not already been acknowledged on the PLB).

On the other hand, the **icbt** instruction can also be used as a convenient mechanism for setting up a fixed, known environment within the instruction cache. This is useful for establishing contents for cache line locking, or for deterministic performance on a particular sequence of code, or even for debugging of low-level hardware and software problems.

When being used for these latter purposes, it is important that the **icbt** instruction deliver a deterministic result, namely the guaranteed establishment in the cache of the specified line. Accordingly, the PPC440x5 core provides a field in the CCR0 register that can be used to cause the **icbt** instruction to operate in this manner. Specifically, when the CCR0 [GICBT] field is set, the execution of **icbt** is *guaranteed* to establish the specified cache line in the instruction cache (assuming that a TLB entry for the referenced memory page exists and has both read and execute permission, and that the caching inhibited storage attribute is not set). The cache line fill associated with such a guaranteed **icbt** will not be abandoned due to subsequent instruction cache misses.

Operation of the **icbt** instruction is affected by the CCR1[FCOM] bit, which forces the **icbt** to appear to miss the cache, even if it should really be a hit. This causes two copies of the line to be established in the cache, simulating a multi-hit parity error. See *Simulating Instruction Cache Parity Errors for Software Testing* on page 111.

#### 4.2.4.5 icread Operation

The **icread** instruction can be used to directly read both the tag and instruction information of a specified word in a specified entry of the instruction cache. The instruction information is read into the Instruction Cache Debug Data Register (ICDBDR), while the tag information is read into a pair of SPRs, the Instruction Cache Debug Tag Register High (ICDBTRH) and Instruction Cache Debug Tag Register Low (ICDBTRL). From there, the information can subsequently be moved into GPRs using **mf spr** instructions.

The execution of the **icread** instruction generates the equivalent of an EA, which is then broken down according to the size of the cache and used to select a specific instruction word from a specific cache line, as shown in *Table 4-4*:

*Table 4-4. Icread and dcread Cache Line Selection*

Array Size	Ignored	Way Selection	Set Selection	Word Selection
8KB	EA <sub>0:18</sub>	EA <sub>19:23</sub>	EA <sub>24:26</sub>	EA <sub>27:29</sub>
16KB	EA <sub>0:17</sub>	EA <sub>18:23</sub>	EA <sub>24:26</sub>	EA <sub>27:29</sub>
32KB	EA <sub>0:16</sub>	EA <sub>17:22</sub>	EA <sub>23:26</sub>	EA <sub>27:29</sub>

The EA generated by the **icread** instruction must be word-aligned (that is, EA<sub>30:31</sub> must be 0); otherwise, it is a programming error and the result is undefined.

If the CCR0[CRPE] bit is set, execution of the **icread** instruction also loads parity information into the ICBDTRH.

Execution of the **icread** instruction is privileged, and is intended for use for debugging purposes only.

#### Programming Note:

The PPC440x5 does not automatically synchronize context between an **icread** instruction and the subsequent **mf spr** instructions which read the results of the **icread** instruction into GPRs. In order to guarantee that the **mf spr** instructions obtain the results of the **icread** instruction, a sequence such as the following must be used:

```

icread    regA,regB    # read cache information (the contents of GPR A and GPR B are
                        # added and the result used to specify a cache line index to be read)
isync
mficbdr   regC         # move instruction information into GPR C
mficbtrh  regD         # move high portion of tag into GPR D
mficbtrl  regE         # move low portion of tag into GPR E
    
```

The following figures illustrate the ICDBDR, ICDBTRH, and ICDBTRL.

0	31
---	----

*Figure 4-7. Instruction Cache Debug Data Register (ICDBDR)*

0:31	Instruction machine code from instruction cache
------	---




Figure 4-8. Instruction Cache Debug Tag Register High (ICDBTRH)

0:23		Tag Effective Address	Bits 0:23 of the 32-bit effective address associated with the cache line read by <b>icread</b> .
24	V	Cache Line Valid 0 Cache line is not valid. 1 Cache line is valid.	The valid indicator for the cache line read by <b>icread</b> .
25:26	TPAR	Tag Parity	The parity bits for the address tag for the cache line read by <b>icread</b> , if CCR0[CRPE] is set.
27	DAPAR	Instruction Data parity	The parity bit for the instruction word at the 32-bit effective address specified in the <b>icread</b> instruction, if CCR0[CRPE] is set.
28:31		Reserved	




Figure 4-9. Instruction Cache Debug Tag Register Low (ICDBTRL)

0:21		Reserved	
22	TS	Translation Space	The address space portion of the virtual address associated with the cache line read by <b>icread</b> .
23	TD	Translation ID (TID) Disable 0 TID enable 1 TID disable	TID Disable field for the memory page associated with the cache line read by <b>icread</b> .
24:31	TID	Translation ID	TID field portion of the virtual address associated with the cache line read by <b>icread</b> .

#### 4.2.4.6 Instruction Cache Parity Operations

The instruction cache contains parity bits and multi-hit detection hardware to protect against soft data errors. Both the instruction tags and data are protected. Instruction cache lines consist of a tag field, 256 bits of data, and 10 parity bits. The tag field is stored in CAM (Content Addressible Memory) cells, while the data and parity bits are stored in normal RAM cells. The instruction cache is real-indexed but virtually-tagged, so the tag field contains a TID field that is compared to the PID value, a TD bit that can be set to disable the TID comparison for shared pages, and the effective address bits to be compared to the fetch request. The exact number of effective address bits depends on the specific cache size.

Two types of errors may be detected by the instruction cache parity logic. In the first type, the parity bits stored in the RAM array are checked against the appropriate data in the instruction cache line when the RAM line is read for an instruction fetch. Note that a parity error will *not* be signaled as a result of an **icread** instruction.

The second type of parity error that may be detected is a multi-hit, sometimes referred to as an MHIT. This type of error may occur when a tag address bit is corrupted, leaving two tags in the instruction cache array that match the same input address. Multi-hit errors may be detected on any instruction fetch. No parity errors of any kind are detected on speculative fetch lookups or **icbt** lookups. Rather, such lookups are treated as cache hits and cause no further action until an instruction fetch lookup at the offending address causes an error to be detected.

If a parity error is detected, and the MSR[ME] is asserted, (i.e. Machine Check interrupts are enabled), the processor vectors to the Machine Check interrupt handler. As is the case for any Machine Check interrupt, after vectoring to the machine check handler, the MCSRR0 contains the value of the oldest "uncommitted" instruction in the pipeline at the time of the exception and MCSRR1 contains the old Machine Status Register (MSR) context. The interrupt handler is able to query Machine Check Status Register (MCSR) to find out that it was called due to a instruction cache parity error, and is then expected to invalidate the I-cache (using **iccci**). The handler returns to the interrupted process using the **rfmci** instruction.

As long as parity checking and machine check interrupts are enabled, instruction cache parity errors are always *recoverable*. That is, they are detected and cause a machine check interrupt before the parity error can cause the machine to update the architectural state with corrupt data. Also note that the machine check interrupt is *asynchronous*; that is, the return address in the MCSRR0 does not point at the instruction address that contains the parity error. Rather, the Machine Check interrupt is taken as soon as the parity error is detected, and some instructions in progress will get flushed and re-executed after the interrupt, just as if the machine were responding to an external interrupt.

#### 4.2.4.7 Simulating Instruction Cache Parity Errors for Software Testing

Because parity errors occur in the cache infrequently and unpredictably, it is desirable to provide users with a way to simulate the effect of an instruction cache parity error so that interrupt handling software may be exercised. This is exactly the purpose of the CCR1[ICDPEI], CCR1[ICTPEI], and CCR1[FCOM] fields.

There are 10 parity bits stored in the RAM cells of each instruction cache line. Two of those bits hold the parity for the tag information, and the remaining 8 bits hold the parity for each of the 8 32-bit instruction words in the line. (There are two parity bits for the tag data because the parity is calculated for alternating bits of the tag field, to guard against a single particle strike event that upsets two adjacent bits. The instruction data bits are physically interleaved in such a way as to allow the use of a single parity bit per instruction word.) The parity bits are calculated and stored as the line is filled into the cache. Usually parity is calculated as the even parity for each set of bits to be protected, which the checking hardware expects. However, if any of the the CCR1[ICTPEI] bits are set, the calculated parity for the corresponding bits of the tag are inverted and stored as odd parity. Similarly, if any of the CCR1[ICDPEI] bits are set, the parity for the corresponding instruction word is set to odd parity. Then, when the instructions stored with odd parity are fetched, they will cause a Parity exception type Machine Check interrupt and exercise the interrupt handling software. The following pseudo-code is an example of how to use the CCR1[ICDPEI] field to simulate a parity error on word 0 of a target cache line:

```

; make sure all this code in the cache before execution
icbi <target line address>    ; get the target line out of the cache
msync                        ; wait for the icbi
mtspr CCR1, 0x80000000      ; Set CCR1[ICDPEI0]
isync                        ; wait for the CCR1 context to update
  
```

```

icbt <target line address> ; this line fills and sets odd parity for word 0
msync                     ; wait for the fill to finish
mtspr CCR1, 0x0          ; Reset CCR1[ICDPEI0]
isync                     ; wait for the CCR1 context to update
br <word 0 of target line> ; fetching the target of the branch causes interrupt

```

Note that any instruction lines filled while bits are set in the CCR1[ICDPEI] or CCR1[ICTPEI] field will be affected, so users must code carefully to affect only the intended addresses.

The CCR1[FCOM] (Force Cache Operation Miss) bit enables the simulation of a multi-hit parity error. When set, it will cause an **icbt** to appear to be a miss, initiating a line fill, *even if the line is really already in the cache*. Thus, this bit allows the same line to be filled to the cache multiple times, which will generate a multi-hit parity error when an attempt is made to fetch an instruction from those cache lines. The following pseudo-code is an example of how to use the CCR1[FCOM] field to simulate a multi-hit parity error in the instruction cache:

```

; make sure all this code is cached and the "target line" is also
; in the cache before execution (use icbt as necessary)
mtspr CCR1, 0x00010000 ; Set CCR1[FCOM]
isync                   ; wait for the CCR1 context to update
icbt <target line address> ; this line fills a second copy of the target line
msync                   ; wait for the fill to finish
mtspr CCR1, 0x0         ; Reset CCR1[FCOM]
isync                   ; wait for the CCR1 context to update
br <word 0 of target line> ; fetching the target of the branch causes interrupt

```

### 4.3 Data Cache Controller

The data cache controller (DCC) handles the execution of the storage access instructions, moving data between memory, the data cache, and the PPC440x5 core GPR file. The DCC interfaces to the PLB using two independent 128-bit interfaces, one for read operations and one for writes. Both of these PLB interfaces support direct attachment to 32-bit and 64-bit PLB subsystems, as well as 128-bit PLB subsystems. The DCC handles frequency synchronization between the PPC440x5 core and the PLB, and can operate at any ratio of  $n:1$ ,  $n:2$ , and  $n:3$ , where  $n$  is an integer greater than the corresponding denominator.

The DCC also handles the execution of the PowerPC data cache management instructions, for touching (prefetching), flushing, invalidating, or zeroing cache lines, or for flash invalidation of the entire cache. Resources for controlling and debugging the data cache operation are also provided.

The DCC interfaces to the Auxiliary Processor (AP) port to provide direct load/store access to the data cache for AP load and store operations, as well as for floating-point load and store instructions. AP load and store instructions can access up to 16 bytes (one quadword) in a single cycle.

Extensive load, store, and flush queues are also provided, such that up to three outstanding line fills, up to four outstanding load misses, and up to two outstanding line flushes can be pending, with the DCC continuing to service subsequent load and store hits in an out-of-order fashion.

The rest of this section describes each of these functions in more detail.



### 4.3.1 DCC Operations

When the DCC executes a load, store, or data cache management instruction, the DCC first translates the effective address specified by the instruction into a real address (see *Memory Management* on page 129 for more information on address translation). Next, the DCC searches the data cache array for the cache line associated with the real address of the requested data. If the cache line is found in the array (a cache *hit*), that cache line is used to satisfy the request, according to the type of operation (load, store, and so on).

If the cache line is *not* found in the array (a cache *miss*), the next action depends upon the type of instruction being executed, as well as the storage attributes of the memory page containing the data being accessed. For most operations, and assuming the memory page is cacheable (see *Caching Inhibited (I)* on page 141), the DCC will send a request for the entire cache line (32 bytes) to the data read PLB interface. The request to the data read PLB interface is sent using the specific byte address requested by the instruction, so that the memory subsystem may read the cache line *target word first* (if it supports such operation) and supply the specific byte[s] requested before retrieving the rest of the cache line.

While the DCC is waiting for a cache line read to complete, it can continue to process subsequent instructions, and handle those accesses that hit in the data cache. That is, the data cache is completely *non-blocking*.

As the DCC receives each portion of the cache line from the data read PLB interface, it is placed into one of three data cache line fill data (DCLFD) buffers. Data from these buffers may be *bypassed* to the GPR file to satisfy load instructions, without waiting for the entire cache line to be filled. Once the entire cache line has been filled into the buffer, it will be written into the data cache at the first opportunity (either when the data cache is otherwise idle, or when subsequent operations require that the DCLFD buffer be written to the data cache).

If a memory subsystem error (such as an address time-out, invalid address, or some other type of hardware error external to the PPC440x5 core) occurs during the filling of the cache line, the line will still be written into the data cache, and data from the line may still be delivered to the GPR file for load instructions. However, the DCC will also report a Data Machine Check exception to the instruction unit of the PPC440x5 core, and a Machine Check interrupt (if enabled) will result. See *Machine Check Interrupt* on page 172 for more information on Machine Check interrupts.

Once a data cache line read request has been made, the entire line read will be performed and the line will be written into the data cache, regardless of whether or not the instruction stream branches (or is interrupted) away from the instruction which prompted the initial line read request. That is, if a data cache line read is initiated speculatively, before knowing whether or not a given instruction execution is really required (for example, on a load instruction which is after an unresolved branch), that line read will be completed, even if it is later determined that the cache line is not really needed. The DCC never aborts any PLB request once it has been made, except when a processor reset occurs while the PLB request is being made.

In general, the DCC will initiate memory read requests without waiting to determine whether the access is actually required by the *sequential execution model (SEM)*. That is, the request will be initiated *speculatively*, even if the instruction causing the request might be abandoned due to a branch, interrupt, or other change in the instruction flow. Of course, write requests to memory cannot be initiated speculatively, although a line fill request in response to a cacheable store access which misses in the data cache could be.

On the other hand, if the guarded storage attribute is set for the memory page being accessed, then the memory request will *not* be initiated until it is guaranteed that the access is required by the SEM. Once initiated, the access will not be abandoned, and the instruction is guaranteed to complete, *prior* to any change in

the instruction stream. That is, if the instruction stream is interrupted, then upon return the instruction execution will resume *after* the instruction which accessed guarded storage, such that the guarded storage access will *not* be re-executed.

See *Guarded (G)* on page 141 for more information on accessing guarded storage.

#### Programming Note:

It is a programming error for a load, store, or **dcbz** instruction to reference a valid cache line in the data cache if the caching inhibited storage attribute is set for the memory page containing the cache line. The result of such an access is undefined. After processor reset, hardware automatically sets the caching inhibited storage attribute for the memory page containing the reset address, and software should flash invalidate the data cache (using **dccci**; see *Data Cache Management and Debug Instruction Summary* on page 121) before executing any load, store, or **dcbz** instructions. Subsequently, lines will not be placed into the data cache unless they are accessed by reference to a memory page for which the caching inhibited attribute has been turned off. If software subsequently turns on the caching inhibited storage attribute for such a page, software must make sure that no lines from that page remain valid in the data cache (typically by using the **dcbf** instruction), before attempting to access the (now caching inhibited) page with load, store, or **dcbz** instructions.

The only instructions that are permitted to reference a caching inhibited line which is a hit in the data cache are the cache management instructions **dcbst**, **dcbf**, **dcbi**, **dccci**, and **dcread**. The **dcbt** and **dcbtst** instructions have no effect if they reference a caching inhibited address, regardless of whether the line exists in the data cache.

#### 4.3.1.1 Load and Store Alignment

The DCC implements all of the integer load and store instructions defined for 32-bit implementations by the PowerPC Book-E architecture. These include byte, halfword, and word loads and stores, as well as load and store string (0 to 127 bytes) and load and store multiple (1 to 32 registers) instructions. Integer byte, halfword, and word loads and stores are performed with a single access to memory if the entire data operand is contained within an aligned 16-byte (quadword) block of memory, regardless of the actual operand alignment within that block. If the data operand crosses a quadword boundary, the load or store is performed using two accesses to memory.

The load and store string and multiple instructions are performed using one memory access for each four bytes, unless and until an access would cross an aligned quadword boundary. The access that would cross the boundary is shortened to access just the number of bytes left within the current quadword block, and then the accesses are resumed with four bytes per access, starting at the beginning of the next quadword block, until the end of the load or store string or multiple is reached.

The DCC handles all misaligned integer load and store accesses in hardware, without causing an Alignment exception. However, the control bit CCR0[FLSTA] can be set to force all misaligned storage access instructions to cause an Alignment exception (see *Figure 4-5* on page 106). When this bit is set, all integer storage accesses must be aligned on an operand-size boundary, or an Alignment exception will result. Load and store multiple instructions must be aligned on a 4-byte boundary, while load and store string instructions can be aligned on any boundary (these instructions are considered to reference *byte* strings, and hence the operand size is a byte).

The DCC also supports load and store operations over the AP interface. These can include floating-point load and store instructions (as defined by PowerPC Book-E), as well as AP load and store instructions for auxiliary processors. While floating-point loads and stores can access either four or eight bytes, AP loads and store can access up to a sixteen bytes.

The DCC handles all misaligned floating-point and AP loads and stores with a single memory access, as long as they do not cross a quadword boundary. If such an access crosses a quadword boundary, the DCC will signal an Alignment exception and an interrupt will result.

The AP interface also supports other options with regards to the handling of misaligned AP and floating-point loads and stores. The AP interface can specify that the DCC should signal an Alignment exception on any AP or floating-point load or store access which is not aligned on either an operand-size boundary or a word boundary. Alternatively, the AP interface can specify that the DCC should *force* the storage access to be aligned on an operand-size boundary by zeroing the appropriate number of low-order address bits.

Floating-point and AP loads and stores are also subject to the function of CCR0[FLSTA].

#### 4.3.1.2 Load Operations

Load instructions that reference cacheable memory pages and miss in the data cache result in cache line read requests being presented to the data read PLB interface. Load operations to caching inhibited memory pages, however, will only access the bytes specifically requested, according to the type of load instruction. This behavior (of only accessing the requested bytes) is only architecturally required when the guarded storage attribute is also set, but the DCC will enforce this requirement on any load to a caching inhibited memory page. Subsequent load operations to the same caching inhibited locations will cause new requests to be sent to the data read PLB interface (data from caching inhibited locations will not be reused from the DCLFD buffer).

The DCC includes three DCLFD buffers, such that a total of three independent data cache line fill requests can be in progress at one time. The DCC can continue to process subsequent load and store accesses while these line fills are in progress.

The DCC also includes a 4-entry *load miss queue (LMQ)*, which holds up to four outstanding load instructions that have either missed in the data cache or access caching inhibited memory pages. Collectively, any LMQ entries which reference cacheable memory pages can reference no more than three different cache lines, since there are only three DCLFD buffers. A load instruction in the LMQ remains there until the requested data arrives in the DCLFD buffer, at which time the data is delivered to the register file and the instruction is removed from the LMQ.

#### 4.3.1.3 Store Operations

The processing of store instructions in the DCC is affected by several factors, including the caching inhibited (I), write-through (W), and guarded (G) storage attributes, as well as whether or not the allocation of data cache lines is enabled for cacheable store misses. There are three different behaviors to consider:

- Whether a data cache line is allocated (if the line is not already in the data cache)
- Whether the data is written directly to memory or only into the data cache
- Whether the store data can be *gathered* with store data from previous or subsequent store instructions before being written to memory

### *Allocation of Data Cache Line on Store Miss*

Of course, if the caching inhibited attribute is set for the memory page being referenced by the store instruction, no data cache line will be allocated. For cacheable store accesses, allocation is controlled by one of two mechanisms: either by a “global” control bit in the Memory Management Unit Control Register (MMUCR), which is applied to all cacheable store accesses regardless of address; or by the U2 storage attribute for the memory page being accessed. See *Memory Management Unit Control Register (MMUCR)* on page 143 for more information on how store miss cache line allocation is controlled.

Regardless of which mechanism is controlling the allocation, if the corresponding bit is set, the cacheable store miss is handled as a *store without allocate (SWOA)*. That is, if SWOA is indicated, then if the access misses in the data cache, then the line will not be allocated (read from memory), and instead the byte[s] being stored will be written directly to memory. Of course, if the cache line has already been allocated and is being read into a DCLFD buffer (due perhaps to a previous cacheable load access), then the SWOA indication is ignored and the access is treated as if it were a store *with allocate*. Similarly, if SWOA is *not* indicated, the cache line *will* be allocated and the cacheable store miss will result in the cache line being read from memory.

### *Direct Write to Memory*

Of course, if the caching inhibited attribute is set for the memory page being referenced by the store instruction, the data must be written directly to memory. For cacheable store accesses that are also write-through, the store data will also be written directly to memory, regardless of whether the access hits in the data cache, and independent of the SWOA mechanism. For cacheable store accesses that are *not* write through, whether the data is written directly to memory depends on both whether the access hits or misses in the data cache, and the SWOA mechanism. If the access is *either* a hit in the data cache, or if SWOA is *not* indicated, then the data will only be written to the data cache, and not to memory. Conversely, if the cacheable store access is both a miss in the data cache and SWOA is indicated, the access will be treated as if it were caching inhibited and the data will be written directly to memory and not to the data cache (since the data cache line is neither there already nor will it be allocated).

### *Store Gathering*

In general, memory write operations caused by separate store instructions that specify locations in either write-through or caching inhibited storage may be gathered into one simultaneous access to memory. Similarly, store accesses that are handled as if they were caching inhibited (due to their being both a miss in the data cache and being indicated as SWOA) may be gathered. Store accesses that are only written into the data cache do not need to be gathered, because there is no performance penalty associated with the separate accesses to the array.

A given sequence of two store operations may only be gathered together if the targeted bytes are contained within the same aligned quadword of memory, and if they are contiguous with respect to each other. Subsequent store operations may continue to be gathered with the previously gathered sequence, subject to the same two rules (same aligned quadword and contiguous with the collection of previously gathered bytes). For example, a sequence of three store word operations to addresses 4, 8, and 0 may all be gathered together, as the first two are contiguous with each other, and the third (store word to address 0) is contiguous with the gathered combination of the previous two.

An additional requirement for store gathering applies to stores which target caching inhibited memory pages. Specifically, a given store to a caching inhibited page can only be gathered with previous store operations if the bytes targeted by the given store do not *overlap* with any of the previously gathered bytes. In other words, a store to a caching inhibited page must be both *contiguous* and *non-overlapping* with the previous store

operation(s) with which it is being gathered. This ensures that the multiple write operations associated with a sequence of store instructions which each target a common caching inhibited location will each be performed independently on that target location.

Finally, a given store operation will *not* be gathered with an earlier store operation if it is separated from the earlier store operation by an **msync** or an **mbar** instruction, or if either of the two store operations reference a memory page which is both guarded and caching inhibited, or if store gathering is disabled altogether by CCR0[DSTG] (see *Figure 4-5* on page 106).

*Table 4-5* summarizes how the various storage attributes and other circumstances affect the DCC behavior on store accesses.

*Table 4-5. Data Cache Behavior on Store Accesses*

Store Access Attributes					DCC Actions		
Caching Inhibited (I)	Hit/Miss	SWOA	Write through (W)	Guarded (G)	Write Cache?	Write Memory?	Gather? <sup>1</sup>
0	Hit	—	0	—	Yes	No	N/A
0	Hit	—	1	0	Yes	Yes	Yes
0	Hit	—	1	1	Yes	Yes	No
0	Miss	0	0	—	Yes	No	N/A
0	Miss	0	1	0	Yes	Yes	Yes
0	Miss	0	1	1	Yes	Yes	No
0	Miss	1	—	0	No	Yes	Yes
0	Miss	1	—	1	No	Yes	No
1	— <sup>2</sup>	—	— <sup>3</sup>	0	No	Yes	Yes <sup>4</sup>
1	— <sup>2</sup>	—	— <sup>3</sup>	1	No	Yes	No

**Note 1:** If store gathering is disabled altogether (by setting CCR0[DSTG] to 1), then such gathering will not occur, regardless of the indication in this table. Furthermore, where this table indicates that store gathering may occur it is presumed that the operations being gathered are targeting the same aligned quadword of memory, and are contiguous with respect to each other.

**Note 2:** It is a programming error for a data cache hit to occur on a store access to a caching inhibited page. The result of such an access is undefined.

**Note 3:** It is programming error for the write-through storage attribute to be set for a page which also has the caching inhibited storage attribute set. The result of an access to such a page is undefined.

**Note 4:** Stores to caching inhibited memory locations may only be gathered with previous store operations if none of the targeted bytes overlap with the bytes targeted by the previous store operations.

#### 4.3.1.4 Line Flush Operations

When a store operation (or the **dcbz** instruction) writes data into the data cache without also writing the data to main memory, the cache line is said to become *dirty*, meaning that the data in the cache is the current value, whereas the value in memory is obsolete. Of course, when such a dirty cache line is replaced (due, for example, to a new cache line fill overwriting the existing line in the cache), the data in the cache line must be copied to memory. Otherwise, the results of the previous store operation[s] that caused the cache line to be marked as dirty would be lost. The operation of copying a dirty cache line to memory is referred to as a *cache line flush*. Cache lines are flushed either due to being replaced when a new cache line is filled, or in response to an explicit software flush request associated with the execution of a **dcbst** or **dcbf** instruction.

The DCC implements four dirty bits per cache line, one for each aligned doubleword within the cache line. Whenever any byte of a given doubleword is stored into a data cache line without also writing that same byte to memory, the corresponding dirty bit for that cache line is set (if CCR1[FFF] is set, then all four dirty bits are set instead of just the one corresponding dirty bit). When a data cache line is flushed, the type of request made to the data write PLB interface depends upon which dirty bits associated with the line are set, and the state of the CCR1[FFF] bit. If the CCR1[FFF] bit is set, the request will always be for an entire 32-byte line. Most users will leave the CCR1[FFF] reset to zero, in which case the controller minimizes the size of the transfer by the following algorithm. If only one dirty bit is set, the request type will be for a single doubleword write. If only two dirty bits are set, and they are in the same quadword, then the request type will be for a 16-byte line write. If two or more dirty bits are set, and they are in different quadwords, the request type will be for an entire 32-byte line write. Regardless of the type of request generated by a cache line flush, the address is always specified as the first byte of the request.

If a store access occurs to a cache line in a memory page for which the write-through storage attribute is set, the dirty bits for that cache line do not get updated, since such a store access will be written directly to memory (and into the data cache as well, if the access is either a hit or if the cache line is allocated upon a miss).

On the other hand, it is permissible for there to exist multiple TLB entries that map to the same real memory page, but specify different values for the write-through storage attribute. In this case, it is possible for a store operation to a virtual page which is marked as non-write-through to have caused the cache line to be marked as dirty, so that a subsequent store operation to a different virtual page mapped to the same real page but marked as write-through encounters a dirty line in the data cache. If this happens, the store to the write-through page will write the data for the store to both the data cache and to memory, but it will not modify the dirty bits for the cache line.

#### 4.3.1.5 Data Read PLB Interface Requests

When a PLB read request results from an access to a cacheable memory location, the request is always for a 32-byte line read, regardless of the type and size of the access that prompted the request. The address presented will be for the first byte of the target of the access.

On the other hand, when a PLB read request results from an access to a caching-inhibited memory location, only the byte[s] specifically accessed will be requested from the PLB, according to the type of instruction prompting the access. Based on the type of storage access instructions (including integer, floating-point, and AP), and based on the mechanism for handling misaligned accesses which cross a quadword boundary (see *Section 4.3.1.1* on page 114), the following types of PLB read requests can occur due to caching inhibited requests:

- 1-byte read (any byte address 0–15 within a quadword)
- 2-byte read (any byte address 0–14 within a quadword)
- 3-byte read (any byte address 0–13 within a quadword)
- 4-byte read (any byte address 0–12 within a quadword)
- 8-byte read (any byte address 0–8 within a quadword)

This request can only occur due to a doubleword floating-point or AP load instruction

- 16-byte line fill (must be for byte address 0 of a quadword)

This request can only occur due to a quadword AP load instruction



#### 4.3.1.6 Data Write PLB Interface Requests

When a PLB write request results from a data cache line flush, the specific type and size of the request is as described in *Line Flush Operations* on page 117.

When a PLB write request results from store operations to caching-inhibited, write-through, and/or store without allocate (SWOA) memory locations, the type and size of the request can be any one of the following (this list includes the possible effects of store gathering; see *Store Gathering* on page 116):

- 1-byte write request (any byte address 0–15 within a quadword)
- 2-byte write request (any byte address 0–14 within a quadword)
- 3-byte write request (any byte address 0–13 within a quadword)
- 4-byte write request (any byte address 0–12 within a quadword)
- 5-byte write request (any byte address 0–11 within a quadword)

Only possible due to store gathering

- 6-byte write request (any byte address 0–10 within a quadword)

Only possible due to store gathering

- 7-byte write request (any byte address 0–9 within a quadword)

Only possible due to store gathering

- 8-byte write request (any byte address 0–8 within a quadword)

Only possible due to store gathering, or due to a floating-point or AP doubleword store

- 9-byte write request (any byte address 0–7 within a quadword)

Only possible due to store gathering

- 10-byte write request (any byte address 0–6 within a quadword)

Only possible due to store gathering

- 11-byte write request (any byte address 0–5 within a quadword)

Only possible due to store gathering

- 12-byte write request (any byte address 0–4 within a quadword)

Only possible due to store gathering

- 13-byte write request (any byte address 0–3 within a quadword)

Only possible due to store gathering

- 14-byte write request (any byte address 0–2 within a quadword)

Only possible due to store gathering

- 15-byte write request (any byte address 0–1 within a quadword)

Only possible due to store gathering

- 16-byte line write request (must be to byte address 0 of a quadword)

Only possible due to store gathering, or due to an AP quadword store

### 4.3.1.7 Storage Access Ordering

In general, the DCC can perform load and store operations *out-of-order* with respect to the instruction stream. That is, the memory accesses associated with a sequence of load and store instructions may be performed in memory in an order different from that implied by the order of the instructions. For example, loads can be processed ahead of earlier stores, or stores can be processed ahead of earlier loads. Also, later loads and stores that hit in the data cache may be processed before earlier loads and stores that miss in the data cache.

The DCC does enforce the requirements of the SEM, such that the net result of a sequence of load and store operations is the same as that implied by the order of the instructions. This means, for example, that if a later load reads the same address written by an earlier store, the DCC guarantees that the load will use the data written by the store, and not the older "pre-store" data. But the memory subsystem could still see a read access associated with an even later load before it sees the write access associated with the earlier store.

If the DCC needs to make a read request to the data read PLB interface, and this request conflicts with (that is, references one or more of the same bytes as) an earlier write request which is being made to the data write PLB interface, the DCC will withhold the read request from the data read PLB interface until the write request has been acknowledged on the data write PLB interface. Once the earlier write request has been acknowledged, the read request will be presented, and it is the responsibility of the PLB subsystem to ensure that the data returned for the read request reflects the value of the data written by the write operation.

Conversely, if a write request conflicts with an earlier read request, the DCC will withhold the write request until the read request has been acknowledged, at which point it is the responsibility of the PLB subsystem to ensure that the data returned for the read request does *not* reflect the newer data being written by the write request.

The PPC440x5 core provides storage synchronization instructions to enable software to control the order in which the memory accesses associated with a sequence of instructions are performed. See *Storage Ordering and Synchronization* on page 81 for more information on the use of these instructions.

### 4.3.2 Data Cache Coherency

The PPC440x5 core does not enforce the coherency of the data cache with respect to alterations of memory performed by entities other than the PPC440x5 core. Similarly, if entities other than the PPC440x5 core attempt to read memory locations which currently exist within the PPC440x5 core data cache and in a modified state, the PPC440x5 core does not recognize such accesses and thus will not respond to such accesses with the modified data from the cache. In other words, the data cache on the PPC440x5 core is not a *snooping* data cache, and there is no hardware enforcement of data cache coherency with memory with respect to other entities in the system which access memory.

It is the responsibility of software to manage this coherency through the appropriate use of the caching inhibited storage attribute, the write-through storage attribute, and/or the data cache management instructions.



### 4.3.3 Data Cache Control and Debug

The PPC440x5 core provides various registers and instructions to control data cache operation and to help debug data cache problems.

#### 4.3.3.1 Data Cache Management and Debug Instruction Summary

For detailed descriptions of the instructions summarized in this section, see *Instruction Set* on page 243

In the instruction descriptions, the term “block” describes the unit of storage operated on by the cache block instructions. For the PPC440x5 core, this is the same as a cache line.

The following instructions are used by software to manage the data cache.

<b>dcba</b>	<p>Data Cache Block Allocate</p> <p>This instruction is implemented as a nop on the PPC440x5 core.</p>
<b>dcbf</b>	<p>Data Cache Block Flush</p> <p>Writes a cache block to memory (if the block has been modified) and then invalidates the block.</p>
<b>dcbi</b>	<p>Data Cache Block Invalidate</p> <p>Invalidates a cache block. Any modified data is discarded and not flushed to memory. Execution of this instruction is privileged.</p>
<b>dcbst</b>	<p>Data Cache Block Store</p> <p>Writes a cache block to memory (if the block has been modified) and leaves the block valid but marked as unmodified.</p>
<b>dcbt</b>	<p>Data Cache Block Touch</p> <p>Initiates a cache block fill, enabling the fill to begin prior to the executing program requiring any data in the block. The program can subsequently access the data in the block without incurring a cache miss.</p>
<b>dcbtst</b>	<p>Data Cache Block Touch for Store</p> <p>Implemented identically to the <b>dcbt</b> instruction.</p>
<b>dcbz</b>	<p>Data Cache Block Set to Zero</p> <p>Establishes a cache line in the data cache and sets the line to all zeros, without first reading the previous contents of the cache block from memory, thereby improving performance. All four doublewords in the line are marked as dirty.</p>
<b>dccci</b>	<p>Data Cache Congruence Class Invalidate</p> <p>Flash invalidates the entire data cache. Execution of this instruction is privileged.</p>
<b>dcread</b>	<p>Data Cache Read</p> <p>Reads a cache line (tag and data) from a specified index of the data cache, into a GPR and a pair of SPRs. Execution of this instruction is privileged.</p> <p>See <i>dcread Operation</i> on page 123.</p>

#### 4.3.3.2 Core Configuration Register 0 (CCR0)

The CCR0 register controls the behavior of the **dcbt** instruction, the handling of misaligned memory accesses, and the store gathering mechanism. The CCR0 register also controls various other functions within the PPC440x5 core that are unrelated to the data cache. Each of these functions is discussed in more detail in the related sections of this manual.

Figure 4-5 on page 106 illustrates the fields of the CCR0 register.

#### 4.3.3.3 Core Configuration Register 1 (CCR1)

The CCR1 register controls the behavior of the line flushes in response to cast-outs or **dcbf** or **dcbst** instructions. It also contains bits to control the artificial injection of parity errors for software testing purposes. Some of those bits affect the data cache, while other control the MMU or instruction cache. Each of these functions is discussed in more detail in the related sections of this manual.

Figure 4-6 on page 107 illustrates the fields of the CCR1 register.

#### 4.3.3.4 *dcbt* and *dcbst* Operation

The **dcbt** instruction is typically used as a “hint” to the processor that a particular block of data is likely to be referenced by the executing program in the near future. Thus the processor can begin filling that block into the data cache, so that when the executing program eventually performs a load from the block it will already be present in the cache, thereby improving performance.

The **dcbst** instruction is typically used for a similar purpose, but specifically for cases where the executing program is likely to *store* to the referenced block in the near future. The differentiation in the purpose of the **dcbst** instruction relative to the **dcbt** instruction is only relevant within shared-memory systems with hardware-enforced support for cache coherency. In such systems, the **dcbst** instruction would attempt to establish the block within the data cache in such a fashion that the processor would most readily be able to subsequently write to the block (for example, in a processor with a *MESI-protocol* cache subsystem, the block might be obtained in *Exclusive* state). However, because the PPC440x5 core does not provide support for hardware-enforced cache coherency, the **dcbst** instruction is handled in an identical fashion to the **dcbt** instruction. The rest of this section thus makes reference only to the **dcbt** instruction, but in all cases the information applies to **dcbst** as well.

Of course, it would not typically be advantageous if the filling of the cache line requested by the **dcbt** itself caused a delay in the reading of data needed by the currently executing program. For this reason, the default behavior of the **dcbt** instruction is for it to be ignored if the filling of the requested cache block cannot be immediately commenced and waiting for such commencement would result in the DCC execution pipeline being stalled. For example, the **dcbt** instruction will be ignored if all three DCLFD buffers are already in use, and execution of subsequent storage access instructions is pending.

On the other hand, the **dcbt** instruction can also be used as a convenient mechanism for setting up a fixed, known environment within the data cache. This is useful for establishing contents for cache line locking, or for deterministic performance on a particular sequence of code, or even for debugging of low-level hardware and software problems.

When being used for these latter purposes, it is important that the **dcbt** instruction deliver a deterministic result, namely the guaranteed establishment in the cache of the specified line. Accordingly, the PPC440x5 core provides a field in the CCR0 register which can be used to cause the **dcbt** instruction to operate in this manner. Specifically, when the CCR0 [GDCBT] field is set, the execution of **dcbt** is *guaranteed* to establish

the specified cache line in the data cache (assuming that a TLB entry for the referenced memory page exists and has read permission, and that the caching inhibited storage attribute is not set). The cache line fill associated with such a guaranteed **dcbt** will occur regardless of any potential instruction execution-stalling circumstances within the DCC.

Operation of the **dcbt** instruction is affected by the CCR1[FCOM] bit, which forces the **dcbt** to appear to miss the cache, even if it should really be a hit. This causes two copies of the line to be established in the cache, simulating a multi-hit parity error. See *Simulating Data Cache Parity Errors for Software Testing* on page 126.

#### 4.3.3.5 *dcread* Operation

The **dcread** instruction can be used to directly read both the tag information and a specified data word in a specified entry of the data cache. The data word is read into the target GPR specified in the instruction encoding, while the tag information is read into a pair of SPRs, Data Cache Debug Tag Register High (DCDBTRH) and Data Cache Debug Tag Register Low (DCDBTRL). The tag information can subsequently be moved into GPRs using **mfspr** instructions.

The execution of the **dcread** instruction generates the equivalent of an EA, which is then broken down according to the size of the cache and used to select a specific data word from a specific cache line, as shown in *Table 4-4* on page 109.

The EA generated by the **dcread** instruction must be word-aligned (that is, EA<sub>30:31</sub> must be 0); otherwise, it is a programming error and the result is undefined.

If the CCR0[CRPE] bit is set, execution of the **dcread** instruction also loads parity information into the DCDBTRL. Note that the DCDBTRL[DPAR] field, unlike all the other parity fields, loads the *check values* of the parity, instead of the raw parity values. That is, the DPAR field will always load with zeros unless a parity error has occurred, or been inserted intentionally using the appropriate bits in the CCR1. This behavior is an artifact of the hardware design of the parity checking logic.

Execution of the **dcread** instruction is privileged, and is intended for use for debugging purposes only.

#### Programming Note:

The PPC440x5 core does not support the use of the **dcread** instruction when the DCC is still in the process of performing cache operations associated with previously executed instructions (such as line fills and line flushes). Also, the PPC440x5 core does not automatically synchronize context between a **dcread** instruction and the subsequent **mfspr** instructions that read the results of the **dcread** instruction into GPRs. In order to guarantee that the **dcread** instruction operates correctly, and that the **mfspr** instructions obtain the results of the **dcread** instruction, a sequence such as the following must be used:

```

msync                # ensure that all previous cache operations have completed
dcread    regT,regA,regB # read cache information; the contents of GPR A and GPR B are
                        # added and the result used to specify a cache line index to be read;
                        # the data word is moved into GPR T and the tag information is read
                        # into DCDBTRH and DCDBTRL
isync                # ensure dcread completes before attempting to read results
mfdcbtrh    regD      # move high portion of tag into GPR D
mfdcbtrl    regE      # move low portion of tag into GPR E
    
```

The following figures illustrate the DCDBTRH and DCDBTRL.




Figure 4-10. Data Cache Debug Tag Register High (DCDBTRH)

0:23	TRA	Tag Real Address	Bits 0:23 of the lower 32 bits of the 36-bit real address associated with the cache line read by <b>dcread</b> .
24	V	Cache Line Valid 0 Cache line is not valid. 1 Cache line is valid.	The valid indicator for the cache line read by <b>dcread</b> .
25:27		Reserved	
28:31	TERA	Tag Extended Real Address	Upper 4 bits of the 36-bit real address associated with the cache line read by <b>dcread</b> .




Figure 4-11. Data Cache Debug Tag Register Low (DCDBTRL)

0:12		Reserved	
13	UPAR	U bit parity	The parity for the U0-U3 bits in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
14:15	TPAR	Tag parity	The parity for the tag bits in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
16:19	DPAR	Data parity	The parity <i>check values</i> for the data bytes in the word read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
20:23	MPAR	Modified (dirty) parity	The parity for the modified (dirty) indicators for each of the four doublewords in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
24:27	D	Dirty Indicators	The "dirty" (modified) indicators for each of the four doublewords in the cache line read by <b>dcread</b> .
28	U0	U0 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
29	U1	U1 Storage Attribute	The U1 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
30	U2	U2 Storage Attribute	The U2 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
31	U3	U3 Storage Attribute	The U3 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .

#### 4.3.3.6 Data Cache Parity Operations

The data cache contains parity bits and multi-hit detection hardware to protect against soft data errors. Both the data cache tags and data are protected. Data cache lines consist of a tag field, 256 bits of data, 4 modified (dirty) bits, 4 user attribute (U) bits, and 39 parity bits. The tag field is stored in CAM (Content Addressible Memory) cells, while the data and parity bits are stored in normal RAM cells. The data cache is physically tagged and indexed, so the tag field contains a real address that is compared to the real address produced by the translation hardware when a load, store, or other cache operation is executed. The exact number of effective address bits depends on the specific cache size.

Two types of errors are detected by the data cache parity logic. In the first type, the parity bits stored in the RAM array are checked against the appropriate data in the RAM line any time the RAM line is read. The RAM data may be read by an indexed operation such as a reload dump (RLD), or by a CAM lookup that matches the tag address, such as a load, **dcbf**, **dcbi**, or **dcbst**. If a line is to be cast out of the cache due to replacement or in response to a **dcbf**, **dcbi**, or **dcbst**, and is determined to have a parity error of this type, no effort is made to prevent the erroneous data from being written onto the PLB. However, the write data on the PLB interface is accompanied by a signal indicating that the data has a parity error.

The second type of parity error that may be detected is a multi-hit, also referred to as an MHIT. This type of error may occur when a tag address bit is corrupted, leaving two tags in the memory array that match the same input. This type of error may be detected on any CAM lookup cycle, such as for stores, loads, **dcbf**, **dcbi**, **dcbst**, **dcbt**, **dcbst**, or **dcbz** instructions. Note that a parity error will *not* be signaled as a result of an **dcread** instruction.

If a parity error is detected and the MSR[ME] is asserted, (i.e. Machine Check interrupts are enabled), the processor vectors to the Machine Check interrupt handler. As is the case for any machine check interrupt, after vectoring to the machine check handler, the MCSRR0 contains the value of the oldest "uncommitted" instruction in the pipeline at the time of the exception and MCSRR1 contains the old (MSR) context. The interrupt handler is able to query Machine Check Status Register (MCSR) to find out that it was called due to a D-cache parity error, and is then expected to either invalidate the data cache (using **dccci**), or to invoke the OS to abort the process or reset the processor, as appropriate. The handler returns to the interrupted process using the **rfmci** instruction.

If the interrupt handler is executed before a parity error is allowed to corrupt the state of the machine, the executing process is *recoverable*, and the interrupt handler can just invalidate the data cache and resume the process. In order to guarantee that all parity errors are recoverable, user code must have two characteristics: first, it must mark all cacheable data pages as "write-through" instead of "copy-back." Second, the software-settable bit (CCR0[PRE]) must be set. This bit forces all load instructions to stall in the last stage of the load/store pipeline for one cycle, but only if needed to ensure that parity errors are recoverable. The pipeline stall guarantees that any parity error is detected and the resulting Machine Check interrupt taken before the load instruction completes and the target GPR is corrupted. Setting CCR0[PRE] degrades overall application performance. However, if the state of the load/store pipeline is such that a load instruction stalls in the last stage for some reason unrelated to parity recoverability, then CCR0[PRE] does not cause an additional cycle stall.

Note that the Parity exception type Machine Check interrupt is *asynchronous*; that is, the return address in the MCSRR0 does not necessarily point at the instruction address that detected the parity error in the data cache. Rather, the Machine Check interrupt is taken as soon as the parity error is detected, and some instructions in progress may get flushed and re-executed after the interrupt, just as if the machine were responding to an external interrupt.

MCSR[DCSP] and MCSR[DCFP] indicate what type of data cache operation caused a parity exception. One of the two bits will be set if a parity error is detected in the data cache, along with MCSR[MCS]. See *Machine Check Interrupts* on page 155.

MCSR[DCSP] is set if a parity error is detected during these search operations:

1. Multi-hit parity errors on any instruction that does a CAM lookup
2. Tag or data parity errors on load instructions
3. Tag parity errors on **dcbf**, **dcbi**, or **dcbst** instructions

MCSR[DCFP] is set if a parity error is detected during these flush operations:

1. Data, dirty, or user parity errors on **dcbf** or **dcbst** instructions
2. Tag, data, dirty, or user parity errors on a line that is cast out for replacement

#### 4.3.3.7 Simulating Data Cache Parity Errors for Software Testing

Because parity errors occur in the cache infrequently and unpredictably, it is desirable to provide users with a way to simulate the effect of a data cache parity error so that interrupt handling software may be exercised. This is exactly the purpose of the CCR1[DCDPEI], CCR1[DCTPEI], CCR1[DCUPEI], CCR1[DCMPEI], and CCR1[FCOM] fields.

The 39 data cache parity bits in each cache line contain one parity bit per data byte (i.e. 32 parity bits per 32 byte line), plus 2 parity bits for the address tag (note that the valid (V) bit, is *not* included in the parity bit calculation for the tag), plus 1 parity bit for the 4-bit U field on the line, plus a parity bit for *each* of the 4 modified (dirty) bits on the line. (There are two parity bits for the tag data because the parity is calculated for alternating bits of the tag field, to guard against a single particle strike event that upsets two adjacent bits. The other data bits are physically interleaved in such a way as to allow the use of a single parity bit per data byte or other field.) All parity bits are calculated and stored as the line is initially filled into the cache. In addition, the data and modified (dirty) parity bits (but not the tag and user parity bits) are updated as the line is updated as the result of executing a store instruction or **dcbz**.

Usually parity is calculated as the even parity for each set of bits to be protected, which the checking hardware expects. However, if any of the the CCR1[DCTPEI] bits are set, the calculated parity for the corresponding bits of the tag are inverted and stored as odd parity. Likewise, if the CCR1[DCUPEI] bit is set, the calculated parity for the user bits is inverted and stored as odd parity. Similarly, if the CCR1[DCDPEI] bit is set, the parity for any data bytes that are written, either during the process of a line fill or by execution of a store instruction, is set to odd parity. Then, when the data stored with odd parity is subsequently loaded, it will cause a Parity exception type Machine Check interrupt and exercise the interrupt handling software. The following pseudocode is an example of how to use the CCR1[DCDPEI] field to simulate a parity error on byte 0 of a target cache line:

```

dcbt <target line address> ; get the target line into the cache
msync                      ; wait for the dcbt
mtspr CCR1, Rx             ; Set CCR1[DCDPEI]
isync                      ; wait for the CCR1 context to update
stb <target byte address> ; store some data at byte 0 of the target line
msync                      ; wait for the store to finish
mtspr CCR1, Rz             ; Reset CCR1[ICDPEI0]
isync                      ; wait for the CCR1 context to update
lb <byte 0 of target line> ; load byte causes interrupt

```

If the CCR1[DCMPEI] bit is set, the parity for any modified (dirty) bits that are written, either during the process of a line fill or by execution of a store instruction or **dcbz**, is set to odd parity. If the CCR1[FFF] bit is also set in addition to CCR1[DCMPEI], then the parity for all four modified (dirty) bits is set to odd parity. Store access to a cache line that is already in the cache and in a memory page for which the write-through storage attribute is set does not update the modified (dirty bits) nor the modified (dirty) parity bits, so for these accesses the CCR1[DCMPEI] setting has no effect.

The CCR1[FCOM] (Force Cache Operation Miss) bit enables the simulation of a multi-hit parity error. When set, it will cause an **dcbt** to appear to be a miss, initiating a line fill, *even if the line is really already in the cache*. Thus, this bit allows the same line to be filled to the cache multiple times, which will generate a multi-hit parity error when an attempt is made to read data from those cache lines. The following pseudocode is an example of how to use the CCR1[FCOM] field to simulate a multi-hit parity error in the data cache:

```

    mtspr CCR0, Rx           ; set CCR0[GDCBT]
    dcbt <target line address> ; this dcbt fills a first copy of the target line, if necessary
    msync                   ; wait for the fill to finish
    mtspr CCR1, Ry          ; set CCR1[FCOM]
    isync                   ; wait for the CCR1 context to update
    dcbt <target line address> ; fill a second copy of the target line
    msync                   ; wait for the fill to finish
    mtspr CCR1, Rz          ; reset CCR1[FCOM]
    isync                   ; wait for the CCR1 context to update
    br <byte 0 of target line> ; load byte causes interrupt
    
```





## 5. Memory Management

The PPC440x5 supports a uniform, 4 gigabyte (GB) effective address (EA) space, and a 64GB (36-bit) real address (RA) space. The PPC440x5 memory management unit (MMU) performs address translation between effective and real addresses, as well as protection functions. With appropriate system software, the MMU supports:

- Translation of effective addresses into real addresses
- Software control of the page replacement strategy
- Page-level access control for instruction and data accesses
- Page-level storage attribute control

### 5.1 MMU Overview

The PPC440x5 generates effective addresses for instruction fetches and data accesses. An effective address is a 32-bit address formed by adding an index or displacement to a base address (see *Effective Address Calculation* on page 41). Instruction effective addresses are for sequential instruction fetches, and for fetches caused by changes in program flow (branches and interrupts). Data effective addresses are for load, store and cache management instructions. The MMU expands effective addresses into virtual addresses (VAs) and then translates them into real addresses (RAs); the instruction and data caches use real addresses to access memory.

The PPC440x5 MMU supports demand-paged virtual memory and other management schemes that depend on precise control of effective to real address mapping and flexible memory protection. Translation misses and protection faults cause precise interrupts. The hardware provides sufficient information to correct the fault and restart the faulting instruction.

The MMU divides storage into pages. The page represents the granularity of address translation, access control, and storage attribute control. PowerPC Book-E architecture defines sixteen page sizes, of which the PPC440x5 MMU supports eight. These eight page sizes (1KB, 4KB, 16KB, 64KB, 256KB, 1MB, 16MB, and 256MB) are simultaneously supported. A valid entry for a page referenced by an effective address must be in the translation lookaside buffer (TLB) in order for the address to be accessed. An attempt to access an address for which no TLB entry exists causes an Instruction or Data TLB Error interrupt, depending on the type of access (instruction or data). See *Interrupts and Exceptions* on page 153 for more information on these and other interrupt types.

The TLB is parity protected against soft errors. If such errors are detected, the CPU can be configured to vector to the machine check interrupt handler, where software can take appropriate action. The details of parity checking and suggested interrupt handling are described below.

#### 5.1.1 Support for PowerPC Book-E MMU Architecture

The Book-E Enhanced PowerPC Architecture defines specific requirements for MMU implementations, but also leaves the details of several features implementation-dependent. The PPC440x5 core is fully compliant with the required MMU mechanisms defined by PowerPC Book-E, but a few optional mechanisms are not supported. These are:

- Memory coherence required (M) storage attribute

Because the PPC440x5 does not provide hardware support for multiprocessor coherence, the memory coherence required storage attribute has no effect. If a TLB entry is created with M=1, then any memory transactions for the page associated with that TLB entry will be indicated as being memory coherence required via a corresponding transfer attribute interface signal, but the setting will have no effect on the operation within the PPC440x5 core.

- TLB Invalidate virtual address (**tlbiva**) instruction

The **tlbiva** instruction is used to support the invalidation of TLB entries in a multiprocessor environment with hardware-enforced coherency, which is not supported by the PPC440x5. Consequently, the attempted execution of this instruction will cause an Illegal Instruction exception type Program interrupt. The **tlbwe** instruction may be used to invalidate TLB entries in a uniprocessor environment.

- TLB Synchronize (**tlbsync**) instruction

The **tlbsync** instruction is used to synchronize software TLB management operations in a multiprocessor environment with hardware-enforced coherency, which is not supported by the PPC440x5. Consequently, this instruction is treated as a no-op.

- Page Sizes

PowerPC Book-E defines sixteen different page sizes, but does not require that an implementation support all of them. Furthermore, some of the page sizes are only applicable to 64-bit implementations, as they are larger than a 32-bit effective address space can support (4GB). Accordingly, the PPC440x5 supports eight of the sixteen page sizes, from 1KB up to 256MB, as mentioned above and as listed in *Table 5-2 Page Size and Effective Address to EPN Comparison* on page 136.

- Address Space

Since the PPC440x5 is a 32-bit implementation of the 64-bit PowerPC Book-E architecture, there are differences in the sizes of some of the TLB fields. First, the Effective Page Number (EPN) field varies from 4 to 22 bits, depending on page size. Second, the page number portion of the real address is made up of a concatenation of two TLB fields, rather than a single Real Page Number (RPN) field as described in PowerPC Book-E. These fields are the RPN field (which can vary from 4 to 22 bits, depending on page size), and the Extended Real Page Number (ERPN) field, which is 4 bits, for a total of 36 bits of real address, when combined with the page offset portion of the real address. See *Address Translation* on page 136 for a more detailed explanation of these fields and the formation of the real address.

## 5.2 Translation Lookaside Buffer

The Translation Lookaside Buffer (TLB) is the hardware resource that controls translation, protection, and storage attributes. A single unified 64-entry, fully-associative TLB is used for both instruction and data accesses. In addition, the PPC440x5 implements two separate, smaller “shadow” TLB arrays, one for instruction fetch accesses and one for data accesses. These shadow TLBs improve performance by lowering the latency for address translation, and by reducing contention for the main unified TLB between instruction fetching and data storage accesses. See *Shadow TLB Arrays* on page 147 for additional information on the operation of the shadow TLB arrays.

Maintenance of TLB entries is under software control. System software determines the TLB entry replacement strategy and the format and use of any page table information. A TLB entry contains all of the information required to identify the page, to specify the address translation, to control the access permissions, and to designate the storage attributes.

A TLB entry is written by copying information from a GPR and the MMUCR[STID] field, using a series of three **tlbwe** instructions. A TLB entry is read by copying the information into a GPR and the MMUCR[STID] field, using a series of three **tlbre** instructions. Software can also search for specific TLB entries using the **tlbsx[.]** instruction. See *TLB Management Instructions* on page 148 for more information on these instructions.

Each TLB entry identifies a page and defines its translation, access controls, and storage attributes. Accordingly, fields in the TLB entry fall into four categories:

- Page identification fields (information required to identify the page to the hardware translation mechanism).
- Address translation fields
- Access control fields
- Storage attribute fields

Table 5-1 summarizes the TLB entry fields for each of the categories.

Table 5-1. TLB Entry Fields

TLB Word	Bit	Field	Description
Page Identification Fields			
0	0:21	EPN	<b>Effective Page Number</b> (variable size, from 4 - 22 bits) Bits 0:n-1 of the EPN field are compared to bits 0:n-1 of the effective address (EA) of the storage access (where n = 32-log <sub>2</sub> (page size in bytes) and page size is specified by the SIZE field of the TLB entry). See Table 5-2 on page 136.
0	22	V	<b>Valid</b> (1 bit) This bit indicates that this TLB entry is valid and may be used for translation. The Valid bit for a given entry can be set or cleared with a <b>tlbwe</b> instruction.
0	23	TS	<b>Translation Address Space</b> (1 bit) This bit indicates the address space this TLB entry is associated with. For instruction fetch accesses, MSR[IS] must match the value of TS in the TLB entry for that TLB entry to provide the translation. Likewise, for data storage accesses (including instruction cache management operations), MSR[DS] must match the value of TS in the TLB entry. For the <b>tlbsx[.]</b> instruction, the MMUCR[STS] field must match the value of TS.
0	24:27	SIZE	<b>Page Size</b> (4 bits) The SIZE field specifies the size of the page associated with the TLB entry as 4 <sup>SIZE</sup> KB, where SIZE ∈ {0, 1, 2, 3, 4, 5, 7, 9}. See Table 5-2 on page 136.
0	28:31	TPAR	<b>Tag Parity</b> (4 bits) The TPAR field reads the parity bits associated with TLB word 0. These bits will be loaded into a GPR as a result of a <b>tlbre</b> , but are ignored when executing a <b>tlbwe</b> , since the parity to be written is calculated by the processor hardware.
0	32:39	TID	<b>Translation ID</b> (8 bits) Field used to identify a globally shared page (TID=0) or the process ID of the owner of a private page (TID<>0). See <i>Page Identification</i> on page 134.

Table 5-1. TLB Entry Fields (continued)

TLB Word	Bit	Field	Description
Address Translation Fields			
1	0:21	RPN	<b>Real Page Number</b> (variable size, from 4 - 22 bits) Bits 0:n-1 of the RPN field are used to replace bits 0:n-1 of the effective address to produce a portion of the real address for the storage access (where $n = 32 - \log_2(\text{page size in bytes})$ and page size is specified by the SIZE field of the TLB entry). Software must set unused low-order RPN bits (that is, bits n:21) to 0. See <i>Address Translation</i> on page 136 and <i>Table 5-3</i> on page 137.
1	22:23	PAR1	<b>Parity for TLB word 1</b> (2 bits) The PAR1 field reads the parity bits associated with TLB word 1. These bits will be loaded into a GPR as a result of a <b>tlbre</b> , but are ignored when executing a <b>tlbwe</b> , since the parity to be written is calculated by the processor hardware.
1	28:31	ERP	<b>Extended Real Page Number</b> (4 bits) The 4-bit ERP field are prepended to the rest of the translated address to form a total of a 36-bit (64 GB) real address. See <i>Address Translation</i> on page 136 and <i>Table 5-3</i> on page 137.
Storage Attribute Fields			
2	0:1	PAR2	<b>Parity for TLB word 2</b> (2 bits) The PAR2 field reads the parity bits associated with TLB word 2. These bits will be loaded into a GPR as a result of a <b>tlbre</b> , but are ignored when executing a <b>tlbwe</b> , since the parity to be written is calculated by the processor hardware..
2	16	U0	<b>User-Definable Storage Attribute 0</b> (1 bit) See <i>User-Definable (U0-U3)</i> on page 142. Specifies the U0 storage attribute for the page associated with the TLB entry. The function of this storage attribute is system-dependent, and has no effect within the PPC440x5 core.
2	17	U1	<b>User-Definable Storage Attribute 1</b> (1 bit) See <i>User-Definable (U0-U3)</i> on page 142. Specifies the U1 storage attribute for the page associated with the TLB entry. The function of this storage attribute is system-dependent, but the PPC440x5 core can be programmed to use this attribute to designate a memory page as containing <i>transient</i> data and/or instructions (see <i>Instruction and Data Caches</i> on page 93).
2	18	U2	<b>User-Definable Storage Attribute 2</b> (1 bit) See <i>User-Definable (U0-U3)</i> on page 142. Specifies the U2 storage attribute for the page associated with the TLB entry. The function of this storage attribute is system-dependent, but the PPC440x5 core can be programmed to use this attribute to specify whether or not stores that miss in the data cache should allocate the line in the data cache (see <i>Instruction and Data Caches</i> on page 93).
2	19	U3	<b>User-Definable Storage Attribute 3</b> (1 bit) See <i>User-Definable (U0-U3)</i> on page 142. Specifies the U3 storage attribute for the page associated with the TLB entry. The function of this storage attribute is system-dependent, and has no effect within the PPC440x5 core.
2	20	W	<b>Write-Through</b> (1 bit) See <i>Write-Through (W)</i> on page 141. 0 The page is not write-through (that is, the page is copy-back). 1 The page is write-through.
2	21	I	<b>Caching Inhibited</b> (1 bit) See <i>Caching Inhibited (I)</i> on page 141. 0 The page is not caching inhibited (that is, the page is cacheable). 1 The page is caching inhibited.

Table 5-1. TLB Entry Fields (continued)

TLB Word	Bit	Field	Description
2	22	M	<b>Memory Coherence Required</b> (1 bit) See <i>Memory Coherence Required (M)</i> on page 141.
			0 The page is not memory coherence required.
			1 The page is memory coherence required.
			Note that the PPC440x5 does not support multiprocessing, and thus all storage accesses will behave as if M=0. Setting M=1 in a TLB entry has no effect other than to cause any system interface transactions to the corresponding page to be indicated as memory coherence required via the "transfer attributes" interface signals.
2	23	G	<b>Guarded</b> (1 bit) See <i>Guarded (G)</i> on page 141.
			0 The page is not guarded.
			1 The page is guarded.
2	24	E	<b>Endian</b> (1 bit) See <i>Endian (E)</i> on page 142.
			0 All accesses to the page are performed with big-endian byte ordering, which means that the byte at the effective address is considered the most-significant byte of a multi-byte scalar (see <i>Byte Ordering</i> on page 42).
			1 All accesses to the page are performed with little-endian byte ordering, which means that the byte at the effective address is considered the least-significant byte of a multi-byte scalar (see <i>Byte Ordering</i> on page 42).
Access Control Fields			
2	26	UX	<b>User State Execute Enable</b> (1 bit) See <i>Execute Access</i> on page 138.
			0 Instruction fetch is not permitted from this page while MSR[PR]=1 and the attempt to execute an instruction from this page while MSR[PR]=1 will cause an Execute Access Control exception type Instruction Storage interrupt.
			1 Instruction fetch and execution is permitted from this page while MSR[PR]=1.
2	27	UW	<b>User State Write Enable</b> (1 bit) See <i>Write Access</i> on page 138.
			0 Store operations and the <b>dcbz</b> instruction are not permitted to this page when MSR[PR]=1 and will cause a Write Access Control exception type Data Storage interrupt.
			1 Store operations and the <b>dcbz</b> instruction are permitted to this page when MSR[PR]=1.
2	28	UR	<b>User State Read Enable</b> (1 bit) See <i>Read Access</i> on page 139.
			0 Load operations and the <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>icbt</b> , and <b>icbi</b> instructions are not permitted from this page when MSR[PR]=1 and will cause a Read Access Control exception. Except for the <b>dcbt</b> , <b>dcbtst</b> , and <b>icbt</b> instructions, a Data Storage interrupt will occur (see <i>Table 5-4</i> on page 140).
			1 Load operations and the <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>icbt</b> , and <b>icbi</b> instructions are permitted from this page when MSR[PR]=1.
2	29	SX	<b>Supervisor State Execute Enable</b> (1 bit) See <i>Execute Access</i> on page 138.
			0 Instruction fetch is not permitted from this page while MSR[PR]=0 and the attempt to execute an instruction from this page while MSR[PR]=0 will cause an Execute Access Control exception type Instruction Storage interrupt.
			1 Instruction fetch and execution is permitted from this page while MSR[PR]=0.
2	30	SW	<b>Supervisor State Write Enable</b> (1 bit) See <i>Write Access</i> on page 138.
			0 Store operations and the <b>dcbz</b> and <b>dcbi</b> instructions are not permitted to this page when MSR[PR]=0 and will cause a Write Access Control exception type Data Storage interrupt.
			1 Store operations and the <b>dcbz</b> and <b>dcbi</b> instructions are permitted to this page when MSR[PR]=0.

Table 5-1. TLB Entry Fields (continued)

TLB Word	Bit	Field	Description
2	31	SR	<b>Supervisor State Read Enable</b> (1 bit) See <i>Read Access</i> on page 139.
			0 Load operations and the <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>icbt</b> , and <b>icbi</b> instructions are not permitted from this page when MSR[PR]=0 and will cause a Read Access Control exception. Except for the <b>dcbt</b> , <b>dcbtst</b> , and <b>icbt</b> instructions, a Data Storage interrupt will occur (see <i>Table 5-4</i> on page 140).
			1 Load operations and the <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>icbt</b> , and <b>icbi</b> instructions are permitted from this page when MSR[PR]=0.

## 5.3 Page Identification

The Valid (V), Effective Page Number (EPN), Translation Space Identifier (TS), Page Size (SIZE), and Translation ID (TID) fields of a particular TLB entry identify the page associated with that TLB entry. Except as noted, all comparisons must succeed to validate this entry for subsequent translation and access control processing. Failure to locate a matching TLB entry based on this criteria for instruction fetches will result in a TLB Miss exception type Instruction TLB Error interrupt. Failure to locate a matching TLB entry based on this criteria for data storage accesses will result in a TLB Miss exception which may result in a Data TLB Error interrupt, depending on the type of data storage access (certain cache management instructions do not result in an interrupt if they cause an exception; they simply no-op).

### 5.3.1 Virtual Address Formation

The first step in page identification is the expansion of the effective address into a virtual address. Again, the effective address is the 32-bit address calculated by a load, store, or cache management instruction, or as part of an instruction fetch. The virtual address is formed by prepending the effective address with a 1-bit address space identifier and an 8-bit process identifier. The process identifier is contained in the Process ID (PID) register. The address space identifier is provided by MSR[IS] for instruction fetches, and by MSR[DS] for data storage accesses and cache management operations, including instruction cache management operations. The resulting 41-bit value forms the virtual address, which is then compared to the virtual addresses contained in the TLB entries.

Note that the **tlbsx[.]** instruction also forms a virtual address, for software controlled search of the TLB. This instruction calculates the effective address in the same manner as a data access instruction, but the process identifier and address space identifier are provided by fields in the MMUCR, rather than by the PID and MSR, respectively (see *TLB Search Instruction (tlbsx[.])* on page 148).

### 5.3.2 Address Space Identifier Convention

The address space identifier differentiates between two distinct virtual address spaces, one generally associated with interrupt-handling and other system-level code and/or data, and the other generally associated with application-level code and/or data.

Typically, user mode programs will run with MSR[IS,DS] both set to 1, allowing access to application-level code and data memory pages. Then, on an interrupt, MSR[IS,DS] are both automatically cleared to 0, so that the interrupt handler code and data areas may be accessed using system-level TLB entries (that is, TLB entries with the TS field = 0). It is also possible that an operating system could set up certain system-level code and data areas (and corresponding TLB entries with the TS field = 1) in the application-level address

space, allowing user mode programs running with MSR[IS,DS] set to 1 to access them (system library routines, for example, which may be shared by multiple user mode and/or supervisor mode programs). System-level code wishing to use these areas would have to first set the corresponding MSR[IS,DS] field in order to use the application-level TLB entries, or there would have to be alternative system-level TLB entries set up.

The net of this is that the notion of application-level code running with MSR[IS,DS] set to 1 and using corresponding TLB entries with the TS=1, and conversely system-level code running with MSR[IS,DS] set to 0 and using corresponding TLB entries with TS=0, is by convention. It is possible to run in user mode with MSR[IS,DS] set to 0, and conversely to run in supervisor mode with MSR[IS,DS] set to 1, with the corresponding TLB entries being used. The only fixed requirement in this regard is the fact that MSR[IS,DS] are cleared on an interrupt, and thus there *must* be a TLB entry for the system-level interrupt handler code with TS=0 in order to be able to fetch and execute the interrupt handler itself. Whether or not other system-level code switches MSR[IS,DS] and creates corresponding system-level TLB entries depends upon the operating system environment.

**Programming Note:** Software must ensure that there is always a valid TLB entry with TS=0 and with supervisor mode execute access permission (SX=1) corresponding to the effective address of the interrupt handlers. Otherwise, an Instruction TLB Error interrupt could result upon the fetch of the interrupt handler for some other interrupt type, and the registers holding the state of the routine which was executing at the time of the original interrupt (SRR0/SRR1) could be corrupted. See *Interrupts and Exceptions* on page 153 for more information.

### 5.3.3 TLB Match Process

This virtual address is used to locate the associated entry in the TLB. The address space identifier, the process identifier, and a portion of the effective address of the storage access are compared to the TS, TID, and EPN fields, respectively, of each TLB entry.

The virtual address matches a TLB entry if:

- The valid (V) field of the TLB entry is 1, and
- The value of the address space identifier is equal to the value of the TS field of the TLB entry, and
- Either the value of the process identifier is equal to the value of the TID field of the TLB entry (private page), or the value of the TID field is 0 (globally shared page), and
- The value of bits 0:n-1 of the effective address is equal to the value of bits 0:n-1 of the EPN field of the TLB entry (where  $n = 32 - \log_2$  (page size in bytes) and page size is specified by the value of the SIZE field of the TLB entry). See *Table 5-2 Page Size and Effective Address to EPN Comparison* on page 136.

A TLB Miss exception occurs if there is no matching entry in the TLB for the page specified by the virtual address (except for the **tlbsx[.]** instruction, which simply returns an undefined value to the GPR file and (for **tlbsx.**) sets CR[CR0]<sub>2</sub> to 0). See *TLB Search Instruction (tlbsx[.])* on page 148.

**Programming Note:** Although it is possible for software to create multiple TLB entries that match the same virtual address, doing so is a programming error and the results are undefined.



Figure 5-1 illustrates the criteria for a virtual address to match a specific TLB entry, while Table 5-2 defines the page sizes associated with each SIZE field value, and the associated comparison of the effective address to the EPN field.

Figure 5-1. Virtual Address to TLB Entry Match Process

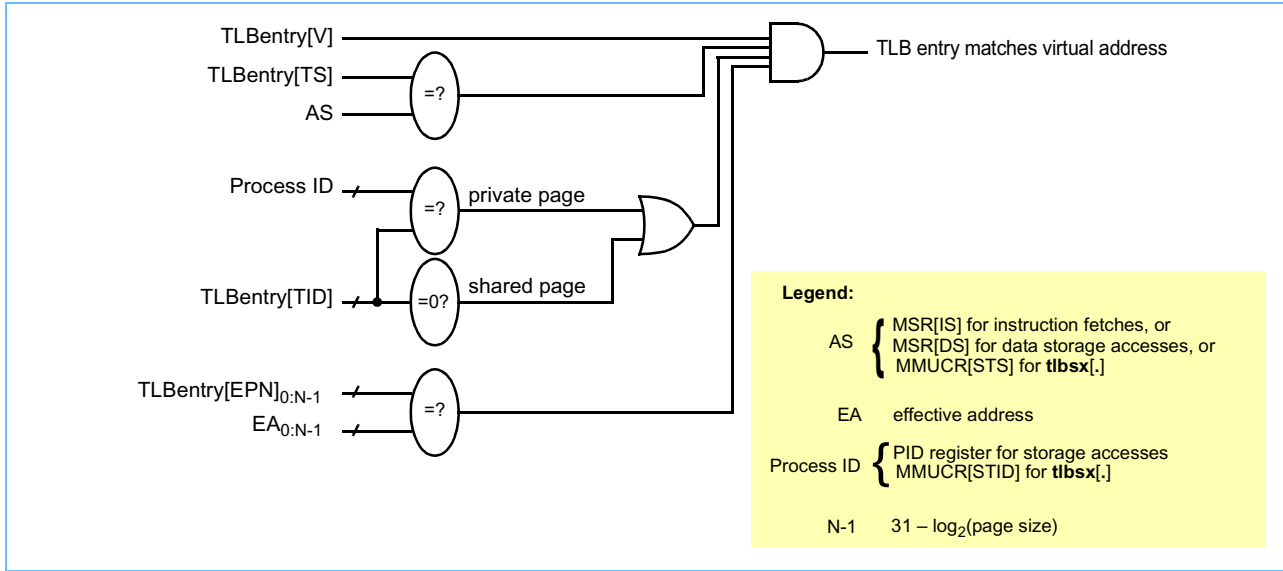


Table 5-2. Page Size and Effective Address to EPN Comparison

SIZE	Page Size	EA to EPN Comparison
0b0000	1KB	EPN <sub>0:21</sub> =? EA <sub>0:21</sub>
0b0001	4KB	EPN <sub>0:19</sub> =? EA <sub>0:19</sub>
0b0010	16KB	EPN <sub>0:17</sub> =? EA <sub>0:17</sub>
0b0011	64KB	EPN <sub>0:15</sub> =? EA <sub>0:15</sub>
0b0100	256KB	EPN <sub>0:13</sub> =? EA <sub>0:13</sub>
0b0101	1MB	EPN <sub>0:11</sub> =? EA <sub>0:11</sub>
0b0110	not supported	not supported
0b0111	16MB	EPN <sub>0:7</sub> =? EA <sub>0:7</sub>
0b1000	not supported	not supported
0b1001	256MB	EPN <sub>0:3</sub> =? EA <sub>0:3</sub>
0b1010	not supported	not supported
0b1011	not supported	not supported
0b1100	not supported	not supported
0b1101	not supported	not supported
0b1110	not supported	not supported
0b1111	not supported	not supported

### 5.4 Address Translation

Once a TLB entry is found which matches the virtual address associated with a given storage access, as described in *Page Identification* on page 134, the virtual address is translated to a real address according to the procedures described in this section.



The Real Page Number (RPN) and Extended Real Page Number (ERP) fields of the matching TLB entry provide the page number portion of the real address. Let  $n=32-\log_2(\text{page size in bytes})$  where *page size* is specified by the SIZE field of the matching TLB entry. Bits  $n:31$  of the effective address (the "page offset") are appended to bits  $0:n-1$  of the RPN field, and bits  $0:3$  of the ERP field are prepended to this value to produce the 36-bit real address (that is,  $RA = ERP_{0:3} \parallel RPN_{0:n-1} \parallel EA_{n:31}$ ).

Figure 5-2 illustrates the address translation process, while Table 5-3 defines the relationship between the different page sizes and the real address formation.

Figure 5-2. Effective-to-Real Address Translation Flow

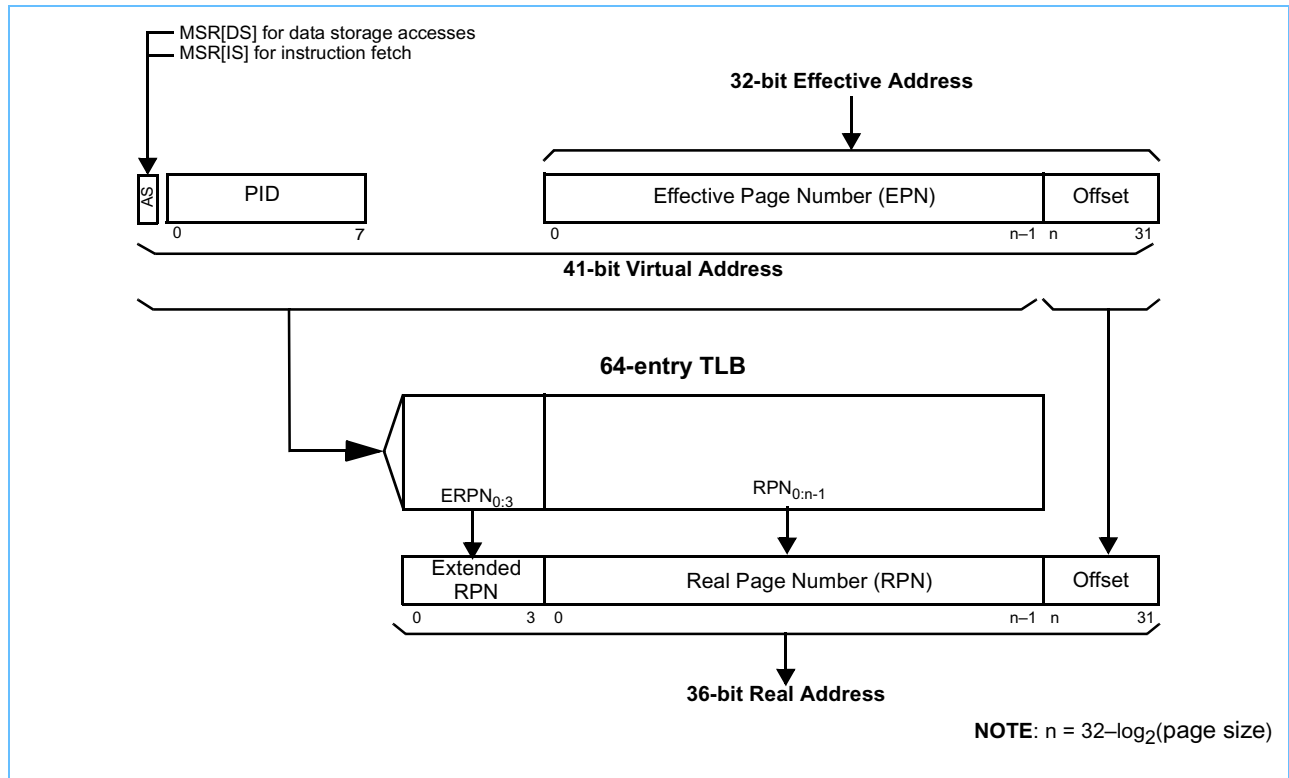


Table 5-3. Page Size and Real Address Formation

SIZE	Page Size	RPN bits required to be 0	Real Address
0b0000	1KB	none	$RPN_{0:21} \parallel EA_{22:31}$
0b0001	4KB	$RPN_{20:21}=0$	$RPN_{0:19} \parallel EA_{20:31}$
0b0010	16KB	$RPN_{18:21}=0$	$RPN_{0:17} \parallel EA_{18:31}$
0b0011	64KB	$RPN_{16:21}=0$	$RPN_{0:15} \parallel EA_{16:31}$
0b0100	256KB	$RPN_{14:21}=0$	$RPN_{0:13} \parallel EA_{14:31}$
0b0101	1MB	$RPN_{12:21}=0$	$RPN_{0:11} \parallel EA_{12:31}$
0b0110	not supported	not supported	not supported
0b0111	16MB	$RPN_{8:21}=0$	$RPN_{0:7} \parallel EA_{8:31}$
0b1000	not supported	not supported	not supported
0b1001	256MB	$RPN_{4:21}=0$	$RPN_{0:3} \parallel EA_{4:31}$
0b1010	not supported	not supported	not supported
0b1011	not supported	not supported	not supported
0b1100	not supported	not supported	not supported
0b1101	not supported	not supported	not supported
0b1110	not supported	not supported	not supported
0b1111	not supported	not supported	not supported

## 5.5 Access Control

Once a matching TLB entry has been identified and the address has been translated, the access control mechanism determines whether the program has execute, read, and/or write access to the page referenced by the address, as described in the following sections.

### 5.5.1 Execute Access

The UX or SX bit of a TLB entry controls *execute* access to a page of storage, depending on the operating mode (user or supervisor) of the processor.

- User mode (MSR[PR] = 1)

Instructions may be fetched and executed from a page in storage while in supervisor mode if the SX access control bit for that page is equal to 1. If the SX access control bit is equal to 0, then instructions from that page will not be fetched, and will not be placed into any cache as the result of a fetch request to that page while in supervisor mode.

Furthermore, if the sequential execution model calls for the execution in supervisor mode of an instruction from a page that is not enabled for execution in supervisor mode (that is, SX=0 when MSR[PR]=0), an Execute Access Control exception type Instruction Storage interrupt is taken (See *Interrupts and Exceptions* on page 153 for more information).

### 5.5.2 Write Access

The UW or SW bit of a TLB entry controls *write* access to a page, depending on the operating mode (user or supervisor) of the processor.

- User mode (MSR[PR] = 1)

Store operations (including the store-class cache management instruction **dcbz**) are permitted to a page in storage while in user mode if the UW access control bit for that page is equal to 1. If execution of a store operation is attempted in user mode to a page for which the UW access control bit is 0, then a Write Access Control exception occurs. If the instruction is an **stswx** with string length 0, then no interrupt is taken and no operation is performed (see *Access Control Applied to Cache Management Instructions* on page 139). For all other store operations, execution of the instruction is suppressed and a Data Storage interrupt is taken.

Note that although the **dcbi** cache management instruction is a store-class instruction, its execution is privileged and thus will not cause a Data Storage interrupt if execution of it is attempted in user mode (a Privileged Instruction exception type Program interrupt will occur instead).

- Supervisor mode (MSR[PR] = 0)

Store operations (including the store-class cache management instructions **dcbz** and **dcbi**) are permitted to a page in storage while in supervisor mode if the SW access control bit for that page is equal to 1. If execution of a store operation is attempted in supervisor mode to a page for which the SW access control bit is 0, then a Write Access Control exception occurs. If the instruction is an **stswx** with string length 0, then no interrupt is taken and no operation is performed (see *Access Control Applied to Cache Management Instructions* on page 139). For all other store operations, execution of the instruction is suppressed and a Data Storage interrupt is taken.

### 5.5.3 Read Access

The UR or SR bit of a TLB entry controls *read* access to a page, depending on the operating mode (user or supervisor) of the processor.

- User mode (MSR[PR] = 1)

Load operations (including the load-class cache management instructions **dcbst**, **dcbf**, **dcbt**, **dcbtst**, **icbi**, and **icbt**) are permitted from a page in storage while in user mode if the UR access control bit for that page is equal to 1. If execution of a load operation is attempted in user mode to a page for which the UR access control bit is 0, then a Read Access Control exception occurs. If the instruction is a load (not including **lswx** with string length 0) or is a **dcbst**, **dcbf**, or **icbi**, then execution of the instruction is suppressed and a Data Storage interrupt is taken. On the other hand, if the instruction is an **lswx** with string length 0, or is a **dcbt**, **dcbtst**, or **icbt**, then no interrupt is taken and no operation is performed (see *Access Control Applied to Cache Management Instructions* below).

- Supervisor mode (MSR[PR] = 0)

Load operations (including the load-class cache management instructions **dcbst**, **dcbf**, **dcbt**, **dcbtst**, **icbi**, and **icbt**) are permitted from a page in storage while in supervisor mode if the SR access control bit for that page is equal to 1. If execution of a load operation is attempted in supervisor mode to a page for which the SR access control bit is 0, then a Read Access Control exception occurs. If the instruction is a load (not including **lswx** with string length 0) or is a **dcbst**, **dcbf**, or **icbi**, then execution of the instruction is suppressed and a Data Storage interrupt is taken. On the other hand, if the instruction is an **lswx** with string length 0, or is a **dcbt**, **dcbtst**, or **icbt**, then no interrupt is taken and no operation is performed (see *Access Control Applied to Cache Management Instructions* below).

### 5.5.4 Access Control Applied to Cache Management Instructions

This section summarizes how each of the cache management instructions is affected by the access control mechanism.

- **dcbz** instructions are treated as *stores* with respect to access control since they actually change the data in a cache block. As such, they can cause Write Access Control exception type Data Storage interrupts.
- **dcbi** instructions are treated as *stores* with respect to access control since they can change the value of a storage location by invalidating the “current” copy of the location in the data cache, effectively “restoring” the value of the location to the “former” value which is contained in memory. As such, they can cause Write Access Control exception type Data Storage interrupts.
- **dcba** instructions are treated as no-ops by the PPC440x5 under all circumstances, and thus can not cause any form of Data Storage interrupt.
- **icbi** instructions are treated as *loads* with respect to access control. As such, they can cause Read Access Control exception type Data Storage interrupts. Note that this instruction may cause a *Data Storage* interrupt (and not an *Instruction Storage* interrupt), even though it otherwise would perform its operation on the *instruction* cache. Instruction storage interrupts are associated with exceptions which occur upon the *fetch* of an instruction, whereas Data storage interrupts are associated with exceptions which occur upon the *execution* of a storage access or cache management instruction.
- **dcbt**, **dcbtst**, and **icbt** instructions are treated as *loads* with respect to access control. As such, they can cause Read Access Control exceptions. However, because these instructions are intended to act merely as “hints” that the specified cache block will likely be accessed by the processor in the near future, such exceptions will not result in a Data Storage interrupt. Instead, if a Read Access Control exception occurs, the instruction is treated as a no-op.

- **dcbf** and **dcbst** instructions are treated as *loads* with respect to access control. As such, they can cause Read Access Control exception type Data Storage interrupts. Flushing or storing a dirty line from the cache is not considered a store since an earlier store operation has already updated the cache line, and the **dcbf** or **dcbst** instruction is simply causing the results of that earlier store operation to be propagated to memory.
- **dccci** and **iccci** instructions do not even generate an address, nor are they affected by the access control mechanism. They are privileged instructions, and if executed in supervisor mode they will flash invalidate the entire associated cache.

Table 5-4 summarizes the effect of access control on each of the cache management instructions.

Table 5-4. Access Control Applied to Cache Management Instructions

Instruction	Read Protection Violation Exception?	Write Protection Violation Exception?
<b>dcba</b>	No	No
<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
<b>dccci</b>	No	No
<b>icbi</b>	Yes	No
<b>icbt</b>	Yes <sup>1</sup>	No
<b>iccci</b>	No	No

1. **dcbt**, **dcbtst**, or **icbt** may cause a Read Access Control exception but will not result in a Data Storage interrupt

## 5.6 Storage Attributes

Each TLB entry specifies a number of storage attributes for the memory page with which it is associated. Storage attributes affect the manner in which storage accesses to a given page are performed. The storage attributes (and their corresponding TLB entry fields) are:

- Write-through (W)
- Caching inhibited (I)
- Memory coherence required (M)
- Guarded (G)
- Endianness (E)
- User-definable (U0, U1, U2, U3)

All combinations of these attributes are supported except combinations which simultaneously specify a region as write-through and caching inhibited.

### 5.6.1 Write-Through (W)

If a memory page is marked as write-through ( $W=1$ ), then the data for all store operations to that page are written to memory, as opposed to only being written into the data cache. If the referenced line also exists in the data cache (that is, the store operation is a “hit”), then the data will also be written into the data cache, although the cache line will *not* be marked as having been modified (that is, the “dirty” bit(s) will not be set).

See *Instruction and Data Caches* on page 93 for more information on the handling of accesses to write-through storage.

### 5.6.2 Caching Inhibited (I)

If a memory page is marked as caching inhibited ( $I=1$ ), then all load, store, and instruction fetch operations perform their access in memory, as opposed to in the respective cache. If  $I=0$ , then the page is cacheable and the operations may be performed in the cache.

It is a programming error for the target location of a load, store, **dcbz**, or fetch access to caching inhibited storage to be in the respective cache; the results of such an access are undefined. It is *not* a programming error for the target locations of the other cache management instructions to be in the cache when the caching inhibited storage attribute is set. The behavior of these instructions is defined for both  $I=0$  and  $I=1$  storage. See the instruction descriptions in *Section 9 Instruction Set* on page 243 for more information.

See *Instruction and Data Caches* on page 93 for more information on the handling of accesses to caching inhibited storage.

### 5.6.3 Memory Coherence Required (M)

The memory coherence required (M) storage attribute is defined by the architecture to support cache and memory coherency within multiprocessor shared memory systems. Because the PPC440x5 does not provide hardware support for multiprocessor coherence, the memory coherence required storage attribute has no effect. If a TLB entry is created with  $M = 1$ , any storage accesses to the page associated with that TLB entry are indicated, using the corresponding transfer attribute interface signal, as being memory coherence required, but the setting has no effect on the operation within the PPC440x5 core.

### 5.6.4 Guarded (G)

The guarded storage attribute is provided to control “speculative” access to “non-well-behaved” memory locations. Storage is said to be “well-behaved” if the corresponding real storage 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. As such, data and instructions can be fetched out-of-order from well-behaved storage without causing undesired side effects.

In general, storage that is not well-behaved should be marked as guarded. Because such storage may represent a control register on an I/O device or may include locations that do not exist, an out-of-order access to such storage may cause an I/O device to perform unintended operations or may result in a Machine Check exception. For example, if the input buffer of a serial I/O device is memory-mapped, then an out-of-order or speculative access to that location could result in the loss of an item of data from the input buffer, if the instruction execution is interrupted and later re-attempted.

A data access to a guarded storage location is performed only if either the access is caused by an instruction that is known to be required by the sequential execution model, or the access is a load and the storage location is already in the data cache. Once a guarded data storage access is initiated, if the storage is also caching inhibited then only the bytes specifically requested are accessed in memory, according to the operand size for the instruction type. Data storage accesses to guarded storage which is marked as cacheable may access the entire cache block, either in the cache itself or in memory.

Instruction fetch is not affected by guarded storage. While the architecture does not prohibit instruction fetching from guarded storage, system software should generally prevent such instruction fetching by marking all guarded pages as “no-execute” (UX/SX = 0). Then, if an instruction fetch is attempted from such a page, the memory access will not occur and an Execute Access Control exception type Instruction Storage interrupt will result if and when execution is attempted for an instruction at any address within the page.

See *Section 4 Instruction and Data Caches* on page 93 for more information on the handling of accesses to guarded storage. Also see *Partially Executed Instructions* on page 158 for information on the relationship between the guarded storage attribute and instruction restart and partially executed instructions.

### 5.6.5 Endian (E)

The endian (E) storage attribute controls the *byte ordering* with which load, store, and fetch operations are performed. Byte ordering refers to the order in which the individual bytes of a multiple-byte scalar operand are arranged in memory. The operands in a memory page with E=0 are arranged with *big-endian* byte ordering, which means that the bytes are arranged with the *most*-significant byte at the lowest-numbered memory address. The operands in a memory page with E=1 are arranged with *little-endian* byte ordering, which means that the bytes are arranged with the *least*-significant byte at the lowest-numbered address.

See *Byte Ordering* on page 42 for a more detailed explanation of big-endian and little-endian byte ordering.

### 5.6.6 User-Definable (U0–U3)

The PPC440x5 core provides four user-definable (U0–U3) storage attributes which can be used to control system-dependent behavior of the storage system. By default, these storage attributes do not have any effect on the operation of the PPC440x5 core, although all storage accesses indicate to the memory subsystem the values of U0–U3 using the corresponding transfer attribute interface signals. The specific system design may then take advantage of these attributes to control some system-level behaviors. As an example, one of the user-definable storage attributes could be used to enable code compression using the IBM CodePack core, if this function is included within a specific implementation incorporating the PPC440x5 core.

On the other hand, the PPC440x5 core can be programmed to make specific use of two of the four user-definable storage attributes. Specifically, by enabling the function using a control bit in the MMUCR (see *Memory Management Unit Control Register (MMUCR)* on page 143), the U1 storage attribute can be used to designate whether storage accesses to the associated memory page should use the “normal” or “transient” region of the respective cache. Similarly, another control bit in the MMUCR can be set to enable the U2 storage attribute to be used to control whether or not store accesses to the associated memory page which miss in the data cache should allocate the line in the cache. The U1 or U2 storage attributes do not affect PPC440x5 core operation unless they are enabled using the MMUCR to perform these specific functions. See *Instruction and Data Caches* on page 93 for more information on the mechanisms that can be controlled by the U1 and U2 storage attributes.

The U0 and U3 storage attributes have no such mechanism that enables them to control any specific function within the PPC440x5 core.

### 5.6.7 Supported Storage Attribute Combinations

Storage modes where both  $W = 1$  and  $I = 1$  (which would represent write-through but caching inhibited storage) are not supported. For all supported combinations of the  $W$  and  $I$  storage attributes, the  $G$ ,  $E$ , and  $U0$ - $U3$  storage attributes may be used in any combination.

## 5.7 Storage Control Registers

In addition to the two registers described below, the  $MSR[IS,DS]$  bits specify which of the two address spaces the respective instruction or data storage accesses are directed towards. Also, the  $MSR[PR]$  bit is used by the access control mechanism. See *Machine State Register (MSR)* on page 159 for more detailed information on the MSR and the function of each of its bits.

### 5.7.1 Memory Management Unit Control Register (MMUCR)

The MMUCR is written from a GPR using `mtspr`, and can be read into a GPR using `mfspir`. In addition, the  $MMUCR[STID]$  is updated with the TID field of the selected TLB entry when a `tlbre` instruction is executed. Conversely, the TID field of the selected TLB entry is updated with the value of the  $MMUCR[STID]$  field when a `tlbwe` instruction is executed. Other functions associated with the STID and other fields of the MMUCR are described in more detail in the sections that follow.

The following figure illustrates the MMUCR.




Figure 5-3. Memory Management Unit Control Register (MMUCR)

0:6		Reserved	
7	SWOA	Store Without Allocate 0 Cacheable store misses allocate a line in the data cache. 1 Cacheable store misses do not allocate a line in the data cache.	If $MMUCR[U2SWOAE] = 1$ , this field is ignored.
8		Reserved	
9	U1TE	U1 Transient Enable 0 Disable U1 storage attribute as transient storage attribute. 1 Enable U1 storage attribute as transient storage attribute.	
10	U2SWOAE	U2 Store without Allocate Enable 0 Disable U2 storage attribute control of store without allocate. 1 Enable U2 storage attribute control of store without allocate.	If $MMUCR[U2SWOAE] = 1$ , the U2 storage attribute overrides $MMUCR[SWOA]$ .
11		Reserved	



12	DULXE	Data Cache Unlock Exception Enable 0 Data cache unlock exception is disabled. 1 Data cache unlock exception is enabled.	<b>dcbf</b> in user mode will cause Cache Locking exception type Data Storage interrupt when MMUCR[DULXE] is 1.
13	IULXE	Instruction Cache Unlock Exception Enable 0 Instruction cache unlock exception is disabled. 1 Instruction cache unlock exception is enabled.	<b>icbi</b> in user mode will cause Cache Locking exception type Data Storage interrupt when MMUCR[IULXE] is 1.
14		Reserved	
15	STS	Search Translation Space	Specifies the value of the translation space (TS) field for the <b>tlbsx[.]</b> instruction
16:23		Reserved	
24:31	STID	Search Translation ID	Specifies the value of the process identifier to be compared against the TLB entry's TID field for the <b>tlbsx[.]</b> instruction; also used to transfer a TLB entry's TID value for the <b>tlbre</b> and <b>tlbwe</b> instructions.

### Store Without Allocate (SWOA) Field

Performance for certain applications can be affected by the allocation of cache lines on store misses. If the store accesses for a particular application are distributed sparsely in memory, and if the data is typically not re-used after having been stored, then performance may be improved by avoiding the latency and bus bandwidth associated with filling the entire cache line containing the bytes being stored. On the other hand, if an application typically stores to contiguous locations, or tends to store repeatedly to the same locations or to re-access data after it has been stored, then performance would likely be improved by allocating the line in the cache upon the first miss so that subsequent accesses will hit in the cache.

The SWOA field is one of two MMUCR fields which can control the allocation of cache lines upon store misses. The other is the U2SWOAE field, and if U2SWOAE is 1 then the U2 storage attribute controls the allocation and the SWOA field is ignored (see *User-Definable (U0–U3)* on page 142). However, if the U2SWOAE field is 0, then the SWOA field controls cache line allocation for all cacheable store misses. Specifically, if a cacheable store access misses in the data cache, then if SWOA is 0, then the cache line will be filled into the data cache, and the store data will be written into the cache (as well as to memory if the associated memory page is also marked as write-through; see *Write-Through (W)* on page 141). Conversely, if SWOA is 1, then cacheable store misses will *not* allocate the line in the data cache, and the store data will be written to memory only, whether or not the write-through attribute is set.

See *Instruction and Data Caches* on page 93 for more information on cache line allocation on store misses.

### U1 Transient Enable (U1TE) Field

When U1TE is 1, then the U1 storage attribute is enabled to control the *transient* mechanism of the instruction and data caches (see *User-Definable (U0–U3)* on page 142). If the U1 field of the TLB entry for the memory page being accessed is 0, then the access will use the *normal* portion of the cache. If the U1 field is 1, then the transient portion of cache will be used.

If the U1TE field is 0, then the transient cache mechanism is disabled and all accesses use the normal portion of the cache.

See Chapter 4, “Instruction and Data Caches” for more information on the transient cache mechanism.



### *U2 Store Without Allocate Enable (U2SWOAE) Field*

An explanation of the allocation of cache lines on store misses is provided in the section on the SWOA field above. The U2SWOAE field is the other mechanism which can control such allocation. If U2SWOAE is 0, then the SWOA field determines whether or not a cache line is allocated on a store miss.

When U2SWOAE is 1, then the U2 storage attribute is enabled to control the allocation on a memory page basis, and the SWOA field is ignored (see *User-Definable (U0–U3)* on page 142). If the U2 field of the TLB entry for the memory page containing the bytes being stored is 0, then the cache line will be allocated in the data cache on a store miss. If the U2 field is 0, then the cache line will *not* be allocated.

See Chapter 4, “Instruction and Data Caches” for more information on cache line allocation on store misses.

### *Data Cache Unlock Exception Enable (DULXE) Field*

The DULXE field can be used to force a Cache Locking exception type Data Storage interrupt to occur if a **dcbf** instruction is executed in user mode (MSR[PR]=1). Since **dcbf** can be executed in user mode and since it causes a cache line to be flushed from the data cache, it has the potential for allowing an application program to remove a locked line from the cache. The locking and unlocking of cache lines is generally a supervisor mode function, as the supervisor has access to the various mechanisms which control the cache locking mechanism (e.g., the Data Cache Victim Limit (DVLIM) and Instruction Cache Victim Limit (IVLIM) registers, and the MMUCR). Therefore, the DULXE field provides a means to prevent any **dcbf** instructions executed while in user mode from flushing any cache lines.

Note that with the PPC440x5 core, the Cache Locking exception occurs independent of whether the target line is truly locked or not. This behavior is necessary because the instruction execution pipeline is such that the exception determination must be made before it is determined whether or not the target line is actually locked (or whether it is even a hit).

Software at the Data Storage interrupt handler can determine whether the target line is locked, and if so whether or not the application should be allowed to unlock it.

If DULXE is 0, or if **dcbf** is executed while in supervisor mode, then the instruction execution is allowed to proceed and flush the target line, independent of whether it is locked or not.

See Chapter 4, “Instruction and Data Caches” for more information on cache locking.

### *Instruction Cache Unlock Exception Enable (IULXE) Field*

The IULXE field can be used to force a Cache Locking exception type Data Storage interrupt to occur if an **icbi** instruction is executed in user mode (MSR[PR]=1). Since **icbi** can be executed in user mode and since it causes a cache line to be removed from the instruction cache, it has the potential for allowing an application program to remove a locked line from the cache. The locking and unlocking of cache lines is generally a supervisor mode function, as the supervisor has access to the various mechanisms which control the cache locking mechanism (e.g., the DVLIM and IVLIM registers, and the MMUCR). Therefore, the IULXE field provides a means to prevent any **icbi** instructions executed while in user mode from flushing any cache lines.

Note that with the PPC440x5 core, the Cache Locking exception occurs independent of whether the target line is truly locked or not. This behavior is necessary because the instruction execution pipeline is such that the exception determination must be made before it is determined whether or not the target line is actually locked (or whether it is even a hit).

Software at the Data Storage interrupt handler can determine whether the target line is locked, and if so whether or not the application should be allowed to unlock it.

If IULXE is 0, or if **icbi** is executed while in supervisor mode, then the instruction execution is allowed to proceed and flush the target line, independent of whether it is locked or not.

See Chapter 4, "Instruction and Data Caches" for more information on cache locking.

#### *Search Translation Space (STS) Field*

The STS field is used by the **tlbsx**[.] instruction to designate the value against which the TS field of the TLB entries is to be matched. For instruction fetch and data storage accesses, the TS field of the TLB entries is compared with the MSR[IS] bit or the MSR[DS] bit, respectively. For **tlbsx**[.] however, the MMUCR[STS] field is used, allowing the TLB to be searched for entries with a TS field which is references an address space other than the one being used by the currently executing process.

See *Address Space Identifier Convention* on page 134 for more information on the TLB entry TS field.

#### *Search Translation ID (STID) Field*

The STID field is used by the **tlbsx**[.] instruction to designate the process identifier value to be compared with the TID field of the TLB entries. For instruction fetch and data storage accesses and cache management operations, the TID field of the TLB entries is compared with the value in the PID register (see *Process ID (PID)* on page 146). For **tlbsx**[.] however, the MMUCR[STID] field is used, allowing the TLB to be searched for entries with a TID field which does not match the Process ID of the currently executing process.

The MMUCR[STID] field is also used to transfer the TLB entry's TID field on **tlbre** and **tlbwe** instructions which target TLB word 0, as there are not enough bits in the GPR used for transferring the other fields such that it could hold this field as well.

See *TLB Match Process* on page 135 for more information on the TLB entry TID field and the address matching process. Also see *TLB Read/Write Instructions (tlbre/tlbwe)* on page 149 for more information on how the MMUCR[STID] field is used by these instructions.

### **5.7.2 Process ID (PID)**

The Process ID (PID) is a 32-bit register, although only the lower 8 bits are defined in the PPC440x5 core. The 8-bit PID value is used as a portion of the virtual address for accessing storage (see *Virtual Address Formation* on page 134). The PID value is compared against the TID field of a TLB entry to determine whether or not the entry corresponds to a given virtual address. If an entry's TID field is 0 (signifying that the entry defines a "global" as opposed to "private" page), then the PID value is ignored when determining whether the entry corresponds to a given virtual address. See *TLB Match Process* on page 135 for a more detailed description of the use of the PID value in the TLB match process.

The PID is written from a GPR using **mtspr**, and can be read into a GPR using **mf spr**. The following figure illustrates the PID.



Figure 5-4. Process ID (PID)

0:23		Reserved
24:31		Process ID

## 5.8 Shadow TLB Arrays

The PPC440x5 core implements two shadow TLB arrays, one for instruction fetches and one for data accesses. These arrays “shadow” the value of a subset of the entries in the main, unified TLB (the UTLB in the context of this discussion). The purpose of the shadow TLB arrays is to reduce the latency of the address translation operation, and to avoid contention for the UTLB array between instruction fetches and data accesses.

The instruction shadow TLB (ITLB) contains four entries, while the data shadow TLB (DTLB) contains eight. There is no latency associated with accessing the shadow TLB arrays, and instruction execution continues in a pipelined fashion as long as the requested address is found in the shadow TLB. If the requested address is not found in the shadow TLB, the instruction fetch or data storage access is automatically stalled while the address is looked up in the UTLB. If the address is found in the UTLB, the penalty associated with the miss in the shadow array is three cycles. If the address is also a miss in the UTLB, then an Instruction or Data TLB Miss exception is reported.

The replacement of entries in the shadow TLB's is managed by hardware, in a round-robin fashion. Upon a shadow TLB miss which leads to a UTLB hit, the hardware will automatically cast-out the oldest entry in the shadow TLB and replace it with the new translation.

The hardware will also automatically invalidate all of the entries in both of the shadow TLB's upon any context synchronization (see *Context Synchronization* on page 80). Context synchronizing operations include the following:

- Any interrupt (including Machine Check)
- Execution of **isync**
- Execution of **rfi**, **r fci**, or **r f mci**
- Execution of **sc**

Note that there are other “context changing” operations which do not cause automatic context synchronization in the hardware. For example, execution of a **tlbwe** instruction changes the UTLB contents but does not cause a context synchronization and thus does not invalidate or otherwise update the shadow TLB entries. In order for changes to the entries in the UTLB (or to other address-related resources such as the PID) to be reflected in the shadow TLB's, software must ensure that a context synchronizing operation occurs prior to any attempt to use any address associated with the updated UTLB entries (either the old or new contents of

those entries). By invalidating the shadow TLB arrays, a context synchronizing operation forces the hardware to refresh the shadow TLB entries with the updated information in the UTLB as each memory page is accessed.

**Note:** Of the items in the preceding list of shadow TLB invalidating operations, the Machine Check interrupt is not architecturally required to be context synchronizing, and thus is not guaranteed to cause invalidation of any shadow TLB arrays on implementations other than the PPC440x5. Consequently, software which is intended to be portable to other implementations should not depend on this behavior, and should insert the appropriate architecturally-defined context synchronizing operation as necessary for desired operation.

## 5.9 TLB Management Instructions

The processor does not imply any format for the page tables or the page table entries. Software has significant flexibility in organizing the size, location, and format of the page table, and in implementing a custom TLB entry replacement strategy. For example, software can “lock” TLB entries that correspond to frequently used storage, so that those entries are never cast out of the TLB, and TLB Miss exceptions to those pages never occur.

In order to enable software to manage the TLB, a set of TLB management instructions is implemented within the PPC440x5 core. These instructions are described briefly in the sections which follow, and in detail in Chapter 9, “Instruction Set.” In addition, the interrupt mechanism provides resources to assist with software handling of TLB-related exceptions. One such resource is Save/Restore Register 0 (SRR0), which provides the exception-causing address for Instruction TLB Error and Instruction Storage interrupts. Another resource is the Data Exception Address Register (DEAR), which provides the exception-causing address for Data TLB Error and Data Storage interrupts. Finally, the Exception Syndrome Register (ESR) provides bits to differentiate amongst the various exception types which may cause a particular interrupt type. See Chapter 6, “Interrupts and Exceptions.” for more information on these mechanisms.

All of the TLB management instructions are privileged, in order to prevent user mode programs from affecting the address translation and access control mechanisms.

### 5.9.1 TLB Search Instruction (**tlbsx**[])

The **tlbsx**[] instruction can be used to locate an entry in the TLB which is associated with a particular virtual address. This instruction forms an effective address for which the TLB is to be searched, in the same manner by which data storage access instructions perform their address calculation, by adding the contents of registers RA (or the value 0 if RA=0) and RB together. The MMUCR[STID] and MMUCR[STS] fields then provide the process ID and address space portions of the virtual address, respectively. Next, the TLB is searched for this virtual address, with the searching process including the notion of disabling the comparison to the process ID if the TID field of a given TLB entry is 0 (see *TLB Match Process* on page 135). Finally, the TLB index of the matching entry is written into the target register (RT). This index value can then serve as the source value for a subsequent **tlbre** or **tlbwe** instruction, to read or update the entry. If no matching entry is found, then the target register contents are undefined.

The “record form” of the instruction (**tlbsx.**) updates CR[CR0]<sub>2</sub> with the result of the search: if a match is found, then CR[CR0]<sub>2</sub> is set to 1; otherwise it is set to 0.

When the TLB is searched using a **tlbsx** instruction, if a matching entry is found, the parity calculated for the tag is compared to the parity stored in the TPAR field. A mismatch causes a parity error exception. Parity errors in words 1 and 2 of the entry will not cause parity error exceptions when executing a **tlbsx** instruction.



## 5.10 Page Reference and Change Status Management

When performing page management, it is useful to know whether a given memory page has been referenced, and whether its contents have been modified. Note that this may be more involved than determining whether a given TLB entry has been used to reference or change memory, since multiple TLB entries may translate to the same memory page. If it is necessary to replace the contents of some memory page with other contents, a page which has been referenced (accessed for any purpose) is more likely to be maintained than a page which has never been referenced. If the contents of a given memory page are to be replaced and the contents of that page have been changed, then the current contents of that page must be written to backup physical storage (such as a hard disk) before replacement.

Similarly, when performing TLB management, it is useful to know whether a given TLB entry has been referenced. When making a decision about which entry of the TLB to replace in order to make room for a new entry, an entry which has never been referenced is a more likely candidate to be replaced.

The PPC440x5 does not automatically record references or changes to a page or TLB entry. Instead, the interrupt mechanism may be used by system software to maintain reference and change information for TLB entries and their associated pages, respectively.

Execute, Read and Write Access Control exceptions may be used to allow software to maintain reference and change information for a TLB entry and for its associated memory page. The following description explains one way in which system software can maintain such reference and change information.

The TLB entry is originally written into the TLB with its access control bits (UX, SX, UR, SR, UW, and SW) off. The first attempt of application code to use the page will therefore cause an Access Control exception and a corresponding Instruction or Data Storage interrupt. The interrupt handler records the reference to the TLB entry and to the associated memory page in a software table, and then turns on the appropriate access control bit, thereby indicating that the particular TLB entry has been referenced. An initial read from the page is handled by only turning on the appropriate UR or SR access control bit, leaving the page "read-only". Subsequent read accesses to the page via that TLB entry will proceed normally.

If a write access is later attempted, a Write Access Control exception type Data Storage interrupt will occur. The interrupt handler records the change status to the memory page in a software table, and then turns on the appropriate UW or SW access control bit, thereby indicating that the memory page associated with the particular TLB entry has been changed. Subsequent write accesses to the page via that TLB entry will proceed normally.

## 5.11 TLB Parity Operations

The TLB is parity protected against soft errors in the TLB memory array that are caused by alpha particle impacts. If such errors are detected, the CPU can be configured to vector to the machine check interrupt handler, which can restore the corrupted state of the TLB from the page tables in system memory.

The TLB is a 64-entry CAM/RAM with 40 tag bits, 41 data bits, and 8 parity bits per entry. Tag and data bits are parity protected with four parity bits for the 40-bit tag, two parity bits for 26 bits of data (i.e. those read and written as word 1 by the **tlbre** and **tlbwe** instructions), and two more parity bits for the remaining 15 bits of data (i.e. word 2). The parity bits are stored in the TLB entries in fields named TPAR, PAR1, and PAR2, respectively. See *Figure 5-5 TLB Entry Word Definitions*



Unlike the instruction and data cache CAM/RAMs, the TLB does *not* detect multiple hits due to parity errors in the tags. The TLB is a relatively small memory array, and the reduction in Soft Error Rate (SER) provided by adding multi-hit detection to the circuit is small, and so, not worth the expense of the feature.

TLB parity bits are *set* any time the TLB is updated, which is always done via a **tlbwe** instruction. TLB parity is *checked* each time the TLB is searched or read, whether to re-fill the ITLB or DTLB, or as a result of a **tlbsx** or **tlbre** instruction. When executing an ITLB or DTLB refill, parity is checked for the tag and both data words. When executing a **tlbsx**, data output is not enabled for the translation and protection outputs of the TLB, so only the tag parity is checked. When executing a **tlbre**, parity is checked only for the word specified in the WS field of the **tlbre** instruction. Detection of a parity error causes a machine check exception. If MSR[ME] is set (which is the usual case), the processor takes a machine check interrupt.

### 5.11.1 Reading TLB Parity Bits with **tlbre**

If CCR0[CRPE] is set, execution of a **tlbre** instruction updates the target register with parity values as well as the tag or other data from the TLB. However, since a **tlbre** that detects a parity error will cause a machine check exception, the target register can only be updated with a “bad” parity value if the MSR[ME] bit is cleared, preventing the machine check interrupt. Thus the usual flow of code that detects a parity error in the TLB and then finds out which entry is erroneous would proceed as:

1. A **tlbre** instruction is executed from normal OS code, resulting in a parity exception. The exception sets MCSR[TLBE] and MCSR[MCS].
2. MSR[ME] = 1, so the CPU vectors to the machine check handler (i.e takes the machine check interrupt) and resets the MSR[ME] bit. Note that even though the parity error causes an *asynchronous* interrupt, that interrupt is guaranteed to be taken before the **tlbre** instruction completes if the CCR0[PRE] (Parity Recoverability Enable) is set, and so the target register (RT) of the **tlbre** will not be updated.
3. The Machine Check handler code includes a series of **tlbre** instructions to query the state of the TLB and find the erroneous entry. When a **tlbre** encounters an erroneous entry and MSR[ME] = 0, the parity exception still happens, setting the MCSR[MCS] and MCSR[TLBE] bits. Additionally, since MSR[ME] = 0, MCSR[IMCE] is set, indicating that an imprecise machine check was detected. Finally, the instruction completes, (since no interrupt is taken because MSR[ME] = 0), updating the target register with data from the TLB, including the parity information.

The **tlbre** causes an exception when it detects a parity error, but the **icread** and **dcread** instructions do not. This inconsistency is explained because OS code commonly uses a sequence of **tlbsx** and **tlbre** instructions to update the “changed” bit in the page table entries. (See section 5.10, “Page Reference and Change Status Management.”) Forcing the software to check the parity manually for each **tlbre** would be a performance limitation. No such functional use exists for the **icread** and **dcread** instructions; they are used only in debugging contexts with no significant performance requirements.

As is the case for any machine check interrupt, after vectoring to the machine check handler, the MCSRR0 contains the value of the oldest “uncommitted” instruction in the pipeline at the time of the exception and MCSRR1 contains the old (MSR) context. The interrupt handler is able to query Machine Check Status Register (MCSR) to find out that it was called due to a TLB parity exception, and then use **tlbre** instructions to find the error in the TLB and restore it from a known good copy in main memory.

**Note:** A parity error on the TLB entry *which maps the machine check exception handler code* prevents recovery. In effect, one of the 64 TLB entries is unprotected, in that the machine cannot recover from an error in that entry. It is possible to add logic to get around this problem, but the reduction in SER achieved by protecting 63 out of 64 TLB entries is sufficient. Further, the software technique of simply dedicating a TLB entry to

the page that contains the machine check handler and periodically refreshing that entry from a known good copy can reduce the probability that the entry will be used with a parity error to near zero.

As mentioned above, any **tlbre** or **tlbsx** instruction that causes a machine check interrupt will be flushed from the pipeline before it completes. Further, any instruction that causes a DTLB or ITLB refill which causes a TLB parity error will be flushed before it completes.

### 5.11.2 Simulating TLB Parity Errors for Software Testing

Because parity errors occur in the TLB infrequently and unpredictably, it is desirable to provide users with a way to simulate the effect of a TLB parity error so that interrupt handling software may be exercised. This is exactly the purpose of the 4-bit CCR1[MMUPEI] field.

Usually, parity is calculated as the even parity for each set of bits to be protected, which the checking hardware expects. This calculation is done as the TLB data is stored with a **tlbwe** instruction. However, if any of the the CCR1[MMUPEI] bits are set, the calculated parity for the corresponding bits of the data being stored are inverted and stored as odd parity. Then, when the data stored with odd parity is subsequently used to refill the DTLB or ITLB, or by a **tlbsx** or **tlbre** instruction, it will cause a Parity exception type Machine Check interrupt and exercise the interrupt handling software. The following pseudo-code is an example of how to use the CCR1[MMUPEI] field to simulate a parity error on a TLB entry:

```

mtspr CCR1, Rx          ; Set some CCR1[MMUPEI] bits
isync                  ; wait for the CCR1 context to update
tlbwe Rs,Ra,0          ; write some data to the TLB with bad parity
tlbwe Rs,Ra,1          ; write some data to the TLB with bad parity
tlbwe Rs,Ra,2          ; write some data to the TLB with bad parity
isync                  ; wait for the tlbwe(s) to finish
mtspr CCR1, Rz          ; Reset CCR1[MMUPEI]
isync                  ; wait for the CCR1 context to update
tlbre RT,RA,WS         ; tlbre with bad parity causes interrupt

```



## 6. Interrupts and Exceptions

This chapter begins by defining the terminology and classification of interrupts and exceptions in *Overview* and *Interrupt Classes*.

*Interrupt Processing* on page 156 explains in general how interrupts are processed, including the requirements for partial execution of instructions.

Several registers support interrupt handling and control. *Interrupt Processing Registers* on page 159 describes these registers.

*Table 6-2 Interrupt and Exception Types* on page 169 lists the interrupts and exceptions handled by the PPC440x5, in the order of Interrupt Vector Offset Register (IVOR) usage. Detailed descriptions of each interrupt type follow, in the same order.

Finally, *Interrupt Ordering and Masking* on page 192 and *Exception Priorities* on page 195 define the priority order for the processing of simultaneous interrupts and exceptions.

### 6.1 Overview

An *interrupt* is the action in which the processor saves its old context (Machine State Register (MSR) and next instruction address) and begins execution at a pre-determined interrupt-handler address, with a modified MSR. *Exceptions* are the events that may cause the processor to take an interrupt, if the corresponding interrupt type is enabled.

Exceptions may be generated by the execution of instructions, or by signals from devices external to the PPC440x5 core, the internal timer facilities, debug events, or error conditions.

### 6.2 Interrupt Classes

All interrupts, except for Machine Check, can be categorized according to two independent characteristics of the interrupt:

- Asynchronous or synchronous
- Critical or non-critical

#### 6.2.1 Asynchronous Interrupts

Asynchronous interrupts are caused by events that are independent of 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.

#### 6.2.2 Synchronous Interrupts

Synchronous interrupts are those that are caused directly by the execution (or attempted execution) of instructions, and are further divided into two classes, *precise* and *imprecise*.

Synchronous, precise interrupts are those that *precisely* indicate the address of the instruction causing the exception that generated the interrupt; or, for certain synchronous, precise interrupt types, the address of the immediately following instruction.

Synchronous, imprecise interrupts are those that may indicate the address of the instruction which caused the exception that generated the interrupt, or the address of some instruction after the one which caused the exception.

### 6.2.2.1 Synchronous, Precise Interrupts

When the execution or attempted execution of an instruction causes a synchronous, precise interrupt, the following conditions exist when the associated interrupt handler begins execution:

- SRR0 (see *Save/Restore Register 0 (SRR0)* on page 161) or CSRR0 (see *Critical Save/Restore Register 0 (CSRR0)* on page 162) addresses either the instruction which caused the exception that generated the interrupt, or the instruction immediately following this instruction. Which instruction is addressed can be determined from a combination of the interrupt type and the setting of certain fields of the ESR (see *Exception Syndrome Register (ESR)* on page 166).
- The interrupt is generated such that all instructions preceding the instruction which caused the exception appear to have completed with respect to the executing processor. However, some storage accesses associated with these preceding instructions may not have been performed with respect to other processors and mechanisms.
- The instruction which caused the exception may appear not to have begun execution (except for having caused the exception), may have been partially executed, or may have completed, depending on the interrupt type (see *Partially Executed Instructions* on page 158).
- Architecturally, no instruction beyond the one which caused the exception has executed.

### 6.2.2.2 Synchronous, Imprecise Interrupts

When the execution or attempted execution of an instruction causes a synchronous, imprecise interrupt, the following conditions exist when the associated interrupt handler begins execution:

- SRR0 or CSRR0 addresses either the instruction which caused the exception that generated the interrupt, or some instruction following this instruction.
- The 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.
- If the imprecise interrupt is forced by the context synchronizing mechanism, due to an instruction that causes another exception that generates an interrupt (for example, Alignment, Data Storage), then SRR0 addresses the interrupt-forcing instruction, and the interrupt-forcing instruction may have been partially executed (see *Partially Executed Instructions* on page 158).
- If the imprecise interrupt is forced by the execution synchronizing mechanism, due to executing an execution synchronizing instruction other than **msync** or **isync**, then SRR0 or CSRR0 addresses the interrupt-forcing instruction, and the interrupt-forcing instruction appears not to have begun execution (except for its forcing the imprecise interrupt). If the imprecise interrupt is forced by an **msync** or **isync** instruction, then SRR0 or CSRR0 may address either the **msync** or **isync** instruction, or the following instruction.
- If the imprecise interrupt is not forced by either the context synchronizing mechanism or the execution synchronizing mechanism, then the instruction addressed by SRR0 or CSRR0 may have been partially executed (see *Partially Executed Instructions* on page 158).

- No instruction following the instruction addressed by SRR0 or CSRR0 has executed.

The only synchronous, imprecise interrupts in the PPC440x5 core are the “special cases” of “delayed” interrupts, which can result when certain kinds of exceptions occur while the corresponding interrupt type is disabled. The first of these is the Floating-Point Enabled exception type Program interrupt. For this type of interrupt to occur, a floating-point unit must be attached to the auxiliary processor interface of the PPC440x5 core, and the Floating-point Enabled Exception Summary bit of the Floating-Point Status and Control Register (FPSCR[FEX]) must be set while Floating-point Enabled exception type Program interrupts are disabled due to MSR[FE0,FE1] both being 0. If and when such interrupts are subsequently enabled, by setting one or the other (or both) of MSR[FE0,FE1] to 1 while FPSCR[FEX] is still 1, then a synchronous, imprecise form of Floating-Point Enabled exception type Program interrupt will occur, and SRR0 will be set to the address of the instruction which would have executed next (that is, the instruction after the one which updated MSR[FE0,FE1]). If the MSR was updated by an **rfi**, **rfci**, or **rfmci** instruction, then SRR0 will be set to the address to which the **rfi**, **rfci**, or **rfmci** was returning, and not to the instruction address which is sequentially after the **rfi**, **rfci**, or **rfmci**.

The second type of delayed interrupt which may be handled as a synchronous, imprecise interrupt is the Debug interrupt. Similar to the Floating-Point Enabled exception type Program interrupt, the Debug interrupt can be temporarily disabled by an MSR bit, MSR[DE]. Accordingly, certain kinds of Debug exceptions may occur and be recorded in the DBSR while MSR[DE] is 0, and later lead to a delayed Debug interrupt if MSR[DE] is set to 1 while a Debug exception is still set in the DBSR. If and when this occurs, the interrupt will either be synchronous and imprecise, or it will be asynchronous, depending on the type of Debug exception causing the interrupt. In either case, CSRR0 is set to the address of the instruction which would have executed next (that is, the instruction after the one which set MSR[DE] to 1). If MSR[DE] is set to 1 by **rfi**, **rfci**, or **rfmci**, then CSRR0 is set to the address to which the **rfi**, **rfci**, or **rfmci** was returning, and not to the address of the instruction which was sequentially after the **rfi**, **rfci**, or **rfmci**.

Besides these special cases of Program and Debug interrupts, all other synchronous interrupts are handled precisely by the PPC440x5 core, including FP Enabled exception type Program interrupts even when the processor is operating in one of the architecturally-defined imprecise modes (MSR[FE0,FE1] = 0b01 or 0b10).

See *Program Interrupt* on page 180 and *Debug Interrupt* on page 188 for a more detailed description of these interrupt types, including both the precise and imprecise cases.

### 6.2.3 Critical and Non-Critical Interrupts

Interrupts can also be classified as critical or noncritical interrupts. Certain interrupt types demand immediate attention, even if other interrupt types are currently being processed and have not yet had the opportunity to save the state of the machine (that is, return address and captured state of the MSR). To enable taking a critical interrupt immediately after a non-critical interrupt has occurred (that is, before the state of the machine has been saved), two sets of Save/Restore Register pairs are provided. Critical interrupts use the Save/Restore Register pair CSRR0/CSRR1. Non-Critical interrupts use Save/Restore Register pair SRR0/SRR1.

### 6.2.4 Machine Check Interrupts

Machine Check interrupts are a special case. They are typically caused by some kind of hardware or storage subsystem failure, or by an attempt to access an invalid address. A Machine Check may be caused indirectly by the execution of an instruction, but not be recognized and/or reported until long after the processor has executed past the instruction that caused the Machine Check. As such, Machine Check interrupts cannot

properly be classified as either synchronous or asynchronous, nor as precise or imprecise. They also do not belong to either the critical or the non-critical interrupt class, but instead have associated with them a unique pair of save/restore registers, Machine Check Save/Restore Registers 0/1 (MCSRR0/1).

Architecturally, the following general rules apply for Machine Check interrupts:

1. No instruction after the one whose address is reported to the Machine Check interrupt handler in MCSRR0 has begun execution.
2. 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 those instructions that are ever going to complete appear to have done so already, and have done so within the context existing prior to the Machine Check interrupt. No further interrupt (other than possible additional Machine Check interrupts) will occur as a result of those instructions.

With the PPC440x5, Machine Check interrupts can be caused by Machine Check exceptions on a memory access for an instruction fetch, for a data access, or for a TLB access. Some of the interrupts generated behave as synchronous, precise interrupts, while other are handled in an asynchronous fashion.

In the case of an Instruction Synchronous Machine Check exception, the PPC440x5 will handle the interrupt as a synchronous, precise interrupt, assuming Machine Check interrupts are enabled ( $MSR[ME] = 1$ ). That is, if a Machine Check exception is detected during an instruction fetch, the exception will not be *reported* to the interrupt mechanism unless and until execution is attempted for the instruction address at which the Machine Check exception occurred. If, for example, the direction of the instruction stream is changed (perhaps due to a branch instruction), such that the instruction at the address associated with the Machine Check exception will not be executed, then the exception will not be reported and no interrupt will occur. If and when an Instruction Machine Check exception is reported, and if Machine Check interrupts are enabled at the time of the reporting of the exception, then the interrupt will be synchronous and precise and MCSRR0 will be set to the instruction address which led to the exception. If Machine Check interrupts are *not* enabled at the time of the reporting of an Instruction Machine Check exception, then a Machine Check interrupt will *not* be generated (*ever*, even if and when  $MSR[ME]$  is subsequently set to 1), although the  $ESR[MCI]$  field will be set to 1 to indicate that the exception has occurred and that the instruction associated with the exception has been executed.

Instruction Asynchronous Machine Check, Data Asynchronous Machine Check, and TLB Asynchronous Machine Check exceptions, on the other hand, are handled in an “asynchronous” fashion. That is, the address reported in MCSRR0 may not be related to the instruction which prompted the access which led, directly or indirectly, to the Machine Check exception. The address may be that of an instruction before or after the exception-causing instruction, or it may reference the exception causing instruction, depending on the nature of the access, the type of error encountered, and the circumstances of the instruction's execution within the processor pipeline. If  $MSR[ME]$  is 0 at the time of a Machine Check exception that is handled in this asynchronous way, a Machine Check interrupt *will* subsequently occur if and when  $MSR[ME]$  is set to 1.

See *Machine Check Interrupt* on page 172 for more detailed information on Machine Check interrupts.

## 6.3 Interrupt Processing

Associated with each kind of interrupt is an *interrupt vector*, that is, the address of the initial instruction that is executed when the corresponding interrupt occurs.

Interrupt processing consists of saving a small part of the processor 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 and the corresponding interrupt type is enabled, the following actions are performed, in order:

1. SRR0 (for non-critical 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 is loaded with information specific to the exception type. Note that many interrupt types can only be caused by a single type of exception, and thus do not need nor use an ESR setting to indicate the cause of the interrupt. Machine Check interrupts load the MCSR.
3. SRR1 (for non-critical class interrupts) or CSRR1 (for critical class interrupts) or MCSRR1 (for Machine Check interrupts) is loaded with a copy of the contents of the MSR.
4. The MSR is updated as described below. The new values take effect beginning with the first instruction following the interrupt.
  - MSR[WE,EE,PR,FP,FE0,DWE,FE1,IS,DS] are set to 0 by all interrupts.
  - MSR[CE,DE] are set to 0 by all critical class interrupts and left unchanged by all non-critical class interrupts.
  - MSR[ME] is set to 0 by Machine Check interrupts and left unchanged by all other interrupts.

See *Machine State Register (MSR)* on page 159 for more detail on the definition of the MSR.

5. Instruction fetching and execution resumes, using the new MSR value, at the interrupt vector address, which is specific to the interrupt type, and is determined as follows:

$$IVPR_{0:15} \parallel IVOR_n_{16:27} \parallel 0b0000$$

where  $n$  specifies the IVOR register to be used for a particular interrupt type (see *Interrupt Vector Offset Registers (IVOR0–IVOR15)* on page 164).

At the end of a non-critical interrupt handling routine, execution of an **rfi** causes the MSR to be restored from the contents of SRR1 and instruction execution to resume at the address contained in SRR0. Likewise, execution of an **rfci** performs the same function at the end of a critical interrupt handling routine, using CSRR0 instead of SRR0 and CSRR1 instead of SRR1. **rfmci** uses MCSRR0 and MCSRR1 in the same manner.

**Programming 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 instructions.

- **stwcx.**, to clear the reservation if one is outstanding, to ensure that a **lwarx** in the “old” process is not paired with a **stwcx.** in the “new” process. See the instruction descriptions for **lwarx** and **stwcx.** in *Instruction Set* on page 243 for more information on storage reservations.
- **msync**, to ensure that all storage operations of an interrupted process are complete with respect to other processors before that process begins executing on another processor.
- **isync**, **rfi**, **rfci**, or **rfmci**, to ensure that the instructions in the “new” process execute in the “new” context.

### 6.3.1 Partially Executed Instructions

In general, the architecture permits load and store instructions to be partially executed, interrupted, and then to be restarted from the beginning upon return from the interrupt. In order to guarantee that a particular load or store instruction will complete without being interrupted and restarted, software must mark the storage being referred to as Guarded, and must use an elementary (not a string or multiple) load or store that is aligned on an operand-sized boundary.

In order to guarantee that load and store instructions can, in general, be restarted and completed correctly without software intervention, the following rules apply when an instruction is partially executed and then interrupted:

- For an elementary load, no part of the target register (GPR(RT), FPR(FRT), or auxiliary processor register) will have been altered.
- For the “update” forms of load and store instructions, the update register, GPR(RA), will not have been altered.

On the other hand, the following effects are permissible when certain instructions are partially executed and then restarted:

- For any store instruction, some of the bytes at the addressed storage location may have been accessed and/or updated (if write access to that page in which bytes were altered is permitted by the access control mechanism). In addition, for the **stwcx** instruction, if the address is not aligned on a word boundary, then the value in CR[CR0] is undefined, as is whether or not the reservation (if one existed) has been cleared.
- For any load, some of the bytes at the addressed storage location may have been accessed (if read access to that page in which bytes were accessed is permitted by the access control mechanism). In addition, for the **lwarx** instruction, if the address is not aligned on a word boundary, it is undefined whether or not a reservation has been set.
- For load multiple and load string instructions, some of the registers in the range to be loaded may have been altered. Including the addressing registers (GPR(RA), and possibly GPR(RB)) in the range to be loaded is an invalid form of these instructions (and a programming error), and thus the rules for partial execution do not protect against overwriting of these registers. Such possible overwriting of the addressing registers makes these invalid forms of load multiple and load strings inherently non-restartable.

In no case will access control be violated.

As previously stated, the only load or store instructions that are guaranteed to not be interrupted after being partially executed are elementary, aligned, guarded loads and stores. All others may 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, guarded):

Critical Input

Machine Check

External Input

Program (Imprecise Mode Floating-Point Enabled)

Note that this type of interrupt can lead to partial execution of a load or store instruction under the architectural definition only; the PPC440x5 core handles the imprecise modes of the Floating-Point Enabled exceptions precisely, and hence this type of interrupt will not lead to partial execution.

Decrementer

Fixed-Interval Timer

Watchdog Timer

Debug (Unconditional Debug Event)

2. Unaligned elementary load or store, or any load or store multiple or string:

All of the above listed under item 1, plus the following:

Alignment

Data Storage (if the access crosses a memory page boundary)

Debug (Data Address Compare, Data Value Compare)

## 6.4 Interrupt Processing Registers

The interrupt processing registers include the Save/Restore Registers (SRR0–SRR1), Critical Save/Restore Registers (CSRR0–CSRR1), Data Exception Address Register (DEAR), Interrupt Vector Offset Registers (IVOR0–IVOR15), Interrupt Vector Prefix Register (IVPR), and Exception Syndrome Register (ESR). Also described in this section is the Machine State Register (MSR), which belongs to the category of processor control registers.

### 6.4.1 Machine State Register (MSR)

The MSR is a register of its own unique type that controls important chip functions, such as the enabling or disabling of various interrupt types.

The MSR can be written from a GPR using the **mtmsr** instruction. The contents of the MSR can be read into a GPR using the **mfmsr** instruction. The MSR[EE] bit can be set or cleared atomically using the **wrtee** or **wrteei** instructions. The MSR contents are also automatically saved, altered, and restored by the interrupt-handling mechanism.

Figure 6-1 shows the MSR bit definitions and describes the function of each bit.




Figure 6-1. Machine State Register (MSR)

0:12		Reserved	
13	WE	Wait State Enable 0 The processor is not in the wait state. 1 The processor is in the wait state.	If MSR[WE] = 1, the processor remains in the wait state until an interrupt is taken, a reset occurs, or an external debug tool clears WE.



14	CE	Critical Interrupt Enable 0 Critical Input and Watchdog Timer interrupts are disabled. 1 Critical Input and Watchdog Timer interrupts are enabled.
15		Reserved
16	EE	External Interrupt Enable 0 External Input, Decrementer, and Fixed Interval Timer interrupts are disabled. 1 External Input, Decrementer, and Fixed Interval Timer interrupts are enabled.
17	PR	Problem State 0 Supervisor state (privileged instructions can be executed) 1 Problem state (privileged instructions can not be executed)
18	FP	Floating Point Available 0 The processor cannot execute floating-point instructions 1 The processor can execute floating-point instructions
19	ME	Machine Check Enable 0 Machine Check interrupts are disabled 1 Machine Check interrupts are enabled.
20	FE0	Floating-point exception mode 0 0 If MSR[FE1] = 0, ignore exceptions mode; if MSR[FE1] = 1, imprecise nonrecoverable mode 1 If MSR[FE1] = 0, imprecise recoverable mode; if MSR[FE1] = 1, precise mode
21	DWE	Debug Wait Enable 0 Disable debug wait mode. 1 Enable debug wait mode.
22	DE	Debug interrupt Enable 0 Debug interrupts are disabled. 1 Debug interrupts are enabled.
23	FE1	Floating-point exception mode 1 0 If MSR[FE0] = 0, ignore exceptions mode; if MSR[FE0] = 1, imprecise recoverable mode 1 If MSR[FE0] = 0, imprecise non-recoverable mode; if MSR[FE0] = 1, precise mode
24:25		Reserved
26	IS	Instruction Address Space 0 All instruction storage accesses are directed to address space 0 (TS = 0 in the relevant TLB entry). 1 All instruction storage accesses are directed to address space 1 (TS = 1 in the relevant TLB entry).
27	DS	Data Address Space 0 All data storage accesses are directed to address space 0 (TS = 0 in the relevant TLB entry). 1 All data storage accesses are directed to address space 1 (TS = 1 in the relevant TLB entry).



28:31		Reserved
-------	--	----------

### 6.4.2 Save/Restore Register 0 (SRR0)

SRR0 is an SPR that is used to save machine state on non-critical interrupts, and to restore machine state when an **rfi** is executed. When a non-critical interrupt occurs, SRR0 is set to an address associated with the process which was executing at the time. When **rfi** is executed, instruction execution returns to the address in SRR0.

In general, SRR0 contains the address of the instruction that caused the non-critical interrupt, or the address of the instruction to return to after a non-critical interrupt is serviced. See the individual descriptions under *Interrupt Definitions* on page 169 for an explanation of the precise address recorded in SRR0 for each non-critical interrupt type.

SRR0 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.

0		29	30	31
---	--	----	----	----

*Figure 6-2. Save/Restore Register 0 (SRR0)*

0:29		Return address for non-critical interrupts
30:31		Reserved

### 6.4.3 Save/Restore Register 1 (SRR1)

SRR1 is an SPR that is used to save machine state on non-critical interrupts, and to restore machine state when an **rfi** is executed. When a non-critical interrupt is taken, the contents of the MSR (prior to the MSR being cleared by the interrupt) are placed into SRR1. When **rfi** is executed, the MSR is restored with the contents of SRR1.

Bits of SRR1 that correspond to reserved bits in the MSR are also reserved.

**Programming Note:** An MSR bit that is reserved may be altered by **rfi**, consistent with the value being restored from SRR1.

SRR1 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 6-3. Save/Restore Register 1 (SRR1)

0:31	Copy of the MSR at the time of a non-critical interrupt.
------	--

### 6.4.4 Critical Save/Restore Register 0 (CSRR0)

CSRR0 is an SPR that is used to save machine state on critical interrupts, and to restore machine state when an **rfci** is executed. When a critical interrupt occurs, CSRR0 is set to an address associated with the process which was executing at the time. When **rfci** is executed, instruction execution returns to 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. See the individual descriptions under *Interrupt Definitions* on page 169 for an explanation of the precise address recorded in CSRR0 for each critical interrupt type.

CSRR0 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 6-4. Critical Save/Restore Register 0 (CSRR0)

0:29	Return address for critical interrupts
30:31	Reserved

### 6.4.5 Critical Save/Restore Register 1 (CSRR1)

CSRR1 is an SPR that is used to save machine state on critical interrupts, and to restore machine state when an **rfci** is executed. When a critical interrupt is taken, the contents of the MSR (prior to the MSR being cleared by the interrupt) are placed into CSRR1. When **rfci** is executed, the MSR is restored with the contents of CSRR1.

Bits of CSRR1 that correspond to reserved bits in the MSR are also reserved.



**Programming Note:** An MSR bit that is reserved may be altered by **rfmci**, consistent with the value being restored from MCSRR1.

MCSRR1 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 6-7. Machine Check Save/Restore Register 1 (MCSRR1)

0:31	Copy of the MSR at the time of a machine check interrupt.
------	---

### 6.4.8 Data Exception Address Register (DEAR)

The 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 exception.

The DEAR can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 6-8. Data Exception Address Register (DEAR)

0:31	Address of data exception for Data Storage, Alignment, and Data TLB Error interrupts
------	--

### 6.4.9 Interrupt Vector Offset Registers (IVOR0–IVOR15)

An IVOR specifies the quadword (16 byte)-aligned interrupt vector offset from the base address provided by the IVPR (see *Interrupt Vector Prefix Register (IVPR)* on page 165) for its respective interrupt type. IVOR0–IVOR15 are provided for the defined interrupt types. The interrupt vector effective address is formed as follows:

$$IVPR_{0:15} \parallel IVOR_n_{16:27} \parallel 0b0000$$

where *n* specifies the IVOR register to be used for the particular interrupt type.

Any IVOR can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.

Figure 6-9 shows the IVOR field definitions, while Table 6-1 identifies the specific IVOR register associated with each interrupt type.




Figure 6-9. Interrupt Vector Offset Registers (IVOR0–IVOR15)

0:15		Reserved
16:27	IVO	Interrupt Vector Offset
28:31		Reserved

Table 6-1. Interrupt Types Associated with each IVOR

IVOR	Interrupt Type
IVOR0	Critical Input
IVOR1	Machine Check
IVOR2	Data Storage
IVOR3	Instruction Storage
IVOR4	External Input
IVOR5	Alignment
IVOR6	Program
IVOR7	Floating Point Unavailable
IVOR8	System Call
IVOR9	Auxiliary Processor Unavailable
IVOR10	Decrementer
IVOR11	Fixed Interval Timer
IVOR12	Watchdog Timer
IVOR13	Data TLB Error
IVOR14	Instruction TLB Error
IVOR15	Debug

#### 6.4.10 Interrupt Vector Prefix Register (IVPR)

The IVPR provides the high-order 16 bits of the effective address of the interrupt vectors, for all interrupt types. The interrupt vector effective address is formed as follows:

$$\text{IVPR}_{0:15} \parallel \text{IVOR}_{n_{16:27}} \parallel 0b0000$$

where  $n$  specifies the IVOR register to be used for the particular interrupt type.

The IVPR can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspr**.




Figure 6-10. Interrupt Vector Prefix Register (IVPR)

0:15	IVP	Interrupt Vector Prefix
16:31		Reserved

### 6.4.11 Exception Syndrome Register (ESR)

The ESR provides a *syndrome* to differentiate between the different kinds of exceptions that can generate the same interrupt type. Upon the generation of one of these types of interrupt, the bit or bits corresponding to the specific exception that generated the interrupt is set, and all other ESR bits are cleared. Other interrupt types do not affect the contents of the ESR. Figure 6-11 provides a summary of the fields of the ESR along with their definitions. See the individual interrupt descriptions under *Interrupt Definitions* on page 169 for an explanation of the ESR settings for each interrupt type, as well as a more detailed explanation of the function of certain ESR fields.

The ESR can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 6-11. Exception Syndrome Register (ESR)

0	MCI	Machine Check—Instruction Fetch Exception 0 Instruction Machine Check exception did not occur. 1 Instruction Machine Check exception occurred.	This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.
1:3		Reserved	
4	PIL	Program Interrupt—Illegal Instruction Exception 0 Illegal Instruction exception did not occur. 1 Illegal Instruction exception occurred.	
5	PPR	Program Interrupt—Privileged Instruction Exception 0 Privileged Instruction exception did not occur. 1 Privileged Instruction exception occurred.	
6	PTR	Program Interrupt—Trap Exception 0 Trap exception did not occur. 1 Trap exception occurred.	

7	FP	<p>Floating Point Operation</p> <p>0 Exception was not caused by a floating point instruction.</p> <p>1 Exception was caused by a floating point instruction.</p>	
8	ST	<p>Store Operation</p> <p>0 Exception was not caused by a store-type storage access or cache management instruction.</p> <p>1 Exception was caused by a store-type storage access or cache management instruction.</p>	
9		Reserved	
10:11	DLK	<p>Data Storage Interrupt—Locking Exception</p> <p>00 Locking exception did not occur.</p> <p>01 Locking exception was caused by <b>dcbf</b>.</p> <p>10 Locking exception was caused by <b>icbi</b>.</p> <p>11 Reserved</p>	
12	AP	<p>AP Operation</p> <p>0 Exception was not caused by an auxiliary processor instruction.</p> <p>1 Exception was caused by an auxiliary processor instruction.</p>	
13	PUO	<p>Program Interrupt—Unimplemented Operation Exception</p> <p>0 Unimplemented Operation exception did not occur.</p> <p>1 Unimplemented Operation exception occurred.</p>	
14	BO	<p>Byte Ordering Exception</p> <p>0 Byte Ordering exception did not occur.</p> <p>1 Byte Ordering exception occurred.</p>	
15	PIE	<p>Program Interrupt—Imprecise Exception</p> <p>0 Exception occurred precisely; SRR0 contains the address of the instruction that caused the exception.</p> <p>1 Exception occurred imprecisely; SRR0 contains the address of an instruction after the one which caused the exception.</p>	<p>This field is only set for a Floating-Point Enabled exception type Program interrupt, and then only when the interrupt occurs imprecisely due to MSR[FE0,FE1] being set to a non-zero value when an attached floating-point unit is already signaling the Floating-Point Enabled exception (that is, FPSCR[FEX] is already 1).</p>
16:26		Reserved	
27	PCRE	<p>Program Interrupt—Condition Register Enable</p> <p>0 Instruction which caused the exception is not a floating-point CR-updating instruction.</p> <p>1 Instruction which caused the exception is a floating-point CR-updating instruction.</p>	<p>This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.</p> <p>This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.</p>
28	PCMP	<p>Program Interrupt—Compare</p> <p>0 Instruction which caused the exception is not a floating-point compare type instruction</p> <p>1 Instruction which caused the exception is a floating-point compare type instruction.</p>	<p>This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.</p> <p>This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.</p>

29:31	PCRF	<p>Program Interrupt—Condition Register Field</p> <p>If ESR[PCRE]=1, this field indicates which CR field was to be updated by the floating-point instruction which caused the exception.</p>	<p>This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.</p> <p>This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.</p>
-------	------	--	--

### 6.4.12 Machine Check Status Register (MCSR)

The MCSR contains status to allow the Machine Check interrupt handler software to determine the cause of a machine check exception. Any Machine Check exception that is handled as an asynchronous interrupt sets MCSR[MCS] and other appropriate bits of the MCSR. If MSR[ME] and MCSR[MCS] are both set, the machine will take a Machine Check interrupt. See *Machine Check Interrupt* on page 172.

The MCSR is read into a GPR using **mfspr**. Clearing the MCSR is performed using **mtspr** by placing a 1 in the GPR source register in all bit positions which are to be cleared in the MCSR, and a 0 in all other bit positions. The data written from the GPR to the MCSR is not direct data, but a mask. A 1 clears the bit and a 0 leaves the corresponding MCSR bit unchanged.




Figure 6-12. Machine Check Status Register (MCSR)

0	MCS	<p>Machine Check Summary</p> <p>0 No async machine check exception pending</p> <p>1 Async machine check exception pending</p>	<p>Set when a machine check exception occurs that is handled in the asynchronous fashion. One of MCSR bits 1:7 will be set simultaneously to indicate the exception type. When MSR[ME] and this bit are both set, Machine Check interrupt is taken.</p>
1	IB	<p>Instruction PLB Error</p> <p>0 Exception not caused by Instruction Read PLB interrupt request (IRQ)</p> <p>1 Exception caused by Instruction Read PLB interrupt request (IRQ)</p>	
2	DRB	<p>Data Read PLB Error</p> <p>0 Exception not caused by Data Read PLB interrupt request (IRQ)</p> <p>1 Exception caused by Data Read PLB interrupt request (IRQ)</p>	
3	DWB	<p>Data Write PLB Error</p> <p>0 Exception not caused by Data Write PLB interrupt request (IRQ)</p> <p>1 Exception caused by Data Write PLB interrupt request (IRQ)</p>	
4	TLBP	<p>Translation Lookaside Buffer Parity Error</p> <p>0 Exception not caused by TLB parity error</p> <p>1 Exception caused by TLB parity error</p>	



5	ICP	Instruction Cache Parity Error 0 Exception not caused by I-cache parity error 1 Exception caused by I-cache parity error	
6	DCSP	Data Cache Search Parity Error 0 Exception not caused by DCU Search parity error 1 Exception caused by DCU Search parity error	Set if and only if the DCU parity error was discovered during a DCU Search operation. See <i>Data Cache Parity Operations</i> on page 125.
7	DCFP	Data Cache Flush Parity Error 0 Exception not caused by DCU Flush parity error 1 Exception caused by DCU Flush parity error	Set if and only if the DCU parity error was discovered during a DCU Flush operation. See <i>Data Cache Parity Operations</i> on page 125.
8	IMPE	Imprecise Machine Check Exception 0 No imprecise machine check exception occurred. 1 Imprecise machine check exception occurred.	Set if a machine check exception occurs that sets MCSR[MCS] (or would if it were not already set) and MSR[ME] = 0.
9:31		Reserved	

## 6.5 Interrupt Definitions

Table 6-2 provides a summary of each interrupt type, in the order corresponding to their associated IVOR register. The table also summarizes the various exception types that may cause that interrupt type; the classification of the interrupt; which ESR bit(s) can be set, if any; and which mask bit(s) can mask the interrupt type, if any.

Detailed descriptions of each of the interrupt types follow the table..

Table 6-2. *Interrupt and Exception Types*

IVOR	Interrupt Type	Exception Type	Asynchronous	Synchronous, Precise	Synchronous, Imprecise	Critical	ESR (See Note 4)	MSR Mask Bit(s)	DBCRO/TCR Mask Bit	Notes
IVOR0	Critical Input	Critical Input	x			x		CE		1
IVOR1	Machine Check	Instruction Machine Check					[MCI]	ME		2
		Data Machine Check						ME		2
		TLB Machine Check						ME		2
IVOR2	Data Storage	Read Access Control		x			[FP,AP]			
		Write Access Control		x			ST,[FP,AP]			
		Cache Locking		x			{DLK <sub>0</sub> ,DLK <sub>1</sub> }			
		Byte Ordering		x			BO,[ST],[FP,AP]			5
IVOR3	Instruction Storage	Execute Access Control		x						
		Byte Ordering		x			BO			6
IVOR4	External Input	External Input	x					EE		1
IVOR5	Alignment	Alignment		x			[ST],[FP,AP]			

Table 6-2. Interrupt and Exception Types (continued)

IVOR	Interrupt Type	Exception Type	Asynchronous	Synchronous, Precise	Synchronous, Imprecise	Critical	ESR (See Note 4)	MSR Mask Bit(s)	DBCR0/TCR Mask Bit	Notes
IVOR6	Program	Illegal Instruction		x			PIL			
		Privileged Instruction		x			PPR,[AP]			
		Trap		x			PTR			
		FP Enabled		x	x		FP,[PIE],[PCRE] {PCMP,PCRF}	FE0 FE1		8
		AP Enabled		x			AP			8
		Unimplemented Op		x			PUO,[FP,AP]			7
IVOR7	FP Unavailable	FP Unavailable		x					8	
IVOR8	System Call	System Call		x						
IVOR9	AP Unavailable	AP Unavailable		x					8	
IVOR10	Decrementer	Decrementer	x				EE	DIE		
IVOR11	Fixed Interval Timer	Fixed Interval Timer	x				EE	FIE		
IVOR12	Watchdog Timer	Watchdog Timer	x			x	CE	WIE		
IVOR13	Data TLB Error	Data TLB Miss		x			[ST],[FP,AP]			
IVOR14	Instruction TLB Error	Instruction TLB Miss		x						
IVOR15	Debug	Trap		x	x	x		DE	IDM	3
		Instruction Address Compare		x	x	x		DE	IDM	3
		Data Address Compare	x	x	x	x		DE	IDM	3
		Data Value Compare	x	x	x	x		DE	IDM	3
		Instruction Complete		x	x	x		DE	IDM	3
		Branch Taken		x		x		DE	IDM	3
		Return		x	x	x		DE	IDM	3
		Interrupt	x			x		DE	IDM	
Unconditional	x			x		DE	IDM			

Table 6-2. Interrupt and Exception Types (continued)

IVOR	Interrupt Type	Exception Type	Asynchronous	Synchronous, Precise	Synchronous, Imprecise	Critical	ESR (See Note 4)	MSR Mask Bit(s)	DBCR0/TCR Mask Bit	Notes
------	----------------	----------------	--------------	----------------------	------------------------	----------	---------------------	-----------------	--------------------	-------

## Table Notes

1. Although it is not specified as part of Book E, it is common for system implementations to provide, as part of the interrupt controller, independent mask and status bits for the various sources of Critical Input and External Input interrupts.
2. Machine Check interrupts are not classified as asynchronous nor synchronous. They are also not classified as critical or non-critical, because they use their own unique set of Save/Restore Registers, MCSRR0/1. See *Machine Check Interrupts* on page 155, and *Machine Check Interrupt* on page 172.
3. Debug exceptions have special rules regarding their interrupt classification (synchronous or asynchronous, and precise or imprecise), depending on the particular debug mode being used and other conditions (see *Debug Interrupt* on page 188).
4. In general, when an interrupt causes a particular ESR bit or bits to be set as indicated in the table, it also causes all other ESR bits to be cleared. Special rules apply to the ESR[MCI] field; see *Machine Check Interrupt* on page 172. If no ESR setting is indicated for any of the exception types within a given interrupt type, then the ESR is unchanged for that interrupt type.

The syntax for the ESR setting indication is as follows:

**[xxx]** means ESR[xxx] *may* be set

**[xxx,yyy,zzz]** means any *one* (or none) of ESR[xxx] or ESR[yyy] or ESR[zzz] *may* be set, but never more than one

**{xxx,yyy,zzz}** means that any combination of ESR[xxx], ESR[yyy], and ESR[zzz] *may* be set, including all or none

**xxx** means ESR[xxx] *will* be set

5. Byte Ordering exception type Data Storage interrupts can only occur when the PPC440x5 core is connected to a floating-point unit or auxiliary processor, and then only when executing FP or AP load or store instructions. See *Data Storage Interrupt* on page 175 for more detailed information on these kinds of exceptions.
6. Byte Ordering exception type Instruction Storage interrupts are defined by the PowerPC Book-E architecture, but cannot occur within the PPC440x5 core. The core is capable of executing instructions from both big endian and little endian code pages.
7. Unimplemented Operation exception type Program interrupts can only occur when the PPC440x5 core is connected to a floating-point unit or auxiliary processor, and then only when executing instruction opcodes which are recognized by the floating-point unit or auxiliary processor but are not implemented within the hardware.
8. Floating-Point Unavailable and Auxiliary Processor Unavailable interrupts, as well as Floating-Point Enabled and Auxiliary Processor Enabled exception type Program interrupts, can only occur when the PPC440x5 core is connected to a floating-point unit or auxiliary processor, and then only when executing instruction opcodes which are recognized by the floating-point unit or auxiliary processor, respectively.

### 6.5.1 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  $MSR[CE] = 1$ . A Critical Input exception is caused by the activation of an asynchronous input to the PPC440x5 core. Although the only mask for this interrupt type within the core is the  $MSR[CE]$  bit, system implementations typically provide an alternative means for independently masking the interrupt requests from the various devices which collectively may activate the PPC440x5 core Critical Input interrupt request input.

**Note:**  $MSR[CE]$  also enables the Watchdog Timer interrupt.

When a Critical Input interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] || IVOR0[IVO] || 0b0000$ .

#### Critical Save/Restore Register 0 (CSRR0)

Set to the effective address of the next instruction to be executed.

#### Critical Save/Restore Register 1 (CSRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

ME: Unchanged.

All other MSR bits set to 0.

**Programming Note:** Software is responsible for taking any action(s) that are required by the implementation in order to clear any Critical Input exception status (such that the Critical Input interrupt request input signal is deasserted) before reenabling  $MSR[CE]$ , in order to avoid another, redundant Critical Input interrupt.

### 6.5.2 Machine Check Interrupt

A Machine Check interrupt occurs when no higher priority exception exists, a Machine Check exception is presented to the interrupt mechanism, and  $MSR[ME] = 1$ . The PowerPC architecture specifies Machine Check interrupts as neither synchronous nor asynchronous, and indeed the exact causes and details of handling such interrupts are implementation dependent. Regardless, for this particular processor core, it is useful to describe the handling of interrupts caused by various types of Machine Check exceptions in those terms. The PPC440x5 core includes four types of Machine Check exceptions. They are:

#### Instruction Synchronous Machine Check exception

A Instruction Synchronous Machine Check exception is caused when timeout or read error is signaled on the instruction read PLB interface during an instruction fetch operation.

Such an exception is not presented to the interrupt handling mechanism, however, unless and until such time as the execution is attempted of an instruction at an address associated with the instruction fetch for which the Instruction Machine Check exception was asserted. When the exception is presented, the  $ESR[MCI]$  bit will be set to indicated the type of exception, regardless of the state of the  $MSR[ME]$  bit.

If  $MSR[ME]$  is 1 when the Instruction Machine Check exception is presented to the interrupt mechanism, then execution of the instruction associated with the exception will be suppressed, a

Machine Check interrupt will occur, and the interrupt processing registers will be updated as described on Page 173. If MSR[ME] is 0, however, then the instruction associated with the exception will be processed as though the exception did not exist and a Machine Check interrupt will *not* occur (*ever*, even if and when MSR[ME] is subsequently set to 1), although the ESR will still be updated as described on Page 173.

### Instruction Asynchronous Machine Check exception

An Instruction Asynchronous Machine Check exception is caused when either:

- an instruction cache parity error is detected
- the read interrupt request is asserted on the instruction read PLB interface.

### Data Asynchronous Machine Check exception

A Data Asynchronous Machine Check exception is caused when one of the following occurs:

- a timeout, read error, or read interrupt request is signaled on the data read PLB interface, during a data read operation
- a timeout, write error, or write interrupt request is signaled on the data write PLB interface, during a data write operation
- a parity error is detected on an access to the data cache.

### TLB Asynchronous Machine Check exception

A TLB Asynchronous Machine Check exception is caused when a parity error is detected on an access to the TLB.

When any Machine Check exception which is handled as an asynchronous interrupt occurs, it is immediately presented to the interrupt handling mechanism. MCSR[MCS] is set, as are other bits of the MCSR as appropriate. A Machine Check interrupt will occur immediately if MSR[ME] is 1, and the interrupt processing registers will be updated as described below. If MSR[ME] is 0, however, then the exception will be “recorded” by the setting of the MCSR[MCS] bit, and deferred until such time as MSR[ME] is subsequently set to 1. Any time the MCSR[MCS] and MSR[ME] are both set to 1, the Machine Check interrupt will be taken. Therefore, MCSR[MCS] must be cleared by software in the Machine Check interrupt handler before executing an **rfmci** to return to processing with MSR[ME] set to 1.

When a Machine Check interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR1[IVO] || 0b0000.

### Machine Check Save/Restore Register 0 (MCSRR0)

For an Instruction Synchronous Machine Check exception, set to the effective address of the instruction presenting the exception. For an Instruction Asynchronous Machine Check, Data Asynchronous Machine Check, or TLB Asynchronous Machine Check exception, set to the effective address of the next instruction to be executed.

### Machine Check Save/Restore Register 1 (MCSRR1)

Set to the contents of the MSR at the time of the interrupt.

### Machine State Register (MSR)

All MSR bits set to 0.

### Exception Syndrome Register (ESR)

MCI Set to 1 for an Instruction Machine Check exception; otherwise left unchanged.

All other defined ESR bits are set to 0 for an Instruction Machine Check exception; otherwise they are left unchanged.

**Programming Note:** If an Instruction Synchronous Machine Check exception is associated with an instruction, and execution of that instruction is attempted while MSR[ME] is 0, then no Machine Check interrupt will occur, but ESR[MCI] will still be set to 1 when the instruction actually executes. Once set, ESR[MCI] cannot be cleared except by software, using the **mtspr** instruction. When processing a Machine Check interrupt handler, software should query ESR[MCI] to determine the type of Machine Check exception, and then clear ESR[MCI]. Then, prior to re-enabling Machine Check interrupts by setting MSR[ME] to 1, software should query the status of ESR[MCI] again to determine whether any additional Instruction Machine Check exceptions have occurred while MSR[ME] was disabled.

### Machine Check Status Register (MCSR)

The MCSR collects status for the Machine Check exceptions that are handled as asynchronous interrupts. MCSR[MCS] is set by any Instruction Asynchronous Machine Check exception, Data Asynchronous Machine Check exception, or TLB Asynchronous Machine Check exception. Other bits in the MCSR are set to indicate the exact type of Machine Check exception.

- MCS Set to 1.
- IB Set to 1 if Instruction Read PLB Interrupt Request (IRQ) is asserted; otherwise set to 0.
- DRB Set to 1 if Data Read PLB Interrupt Request (IRQ) is asserted; otherwise set to 0.
- DWB Set to 1 if Data Write PLB Interrupt Request (IRQ) is asserted; otherwise set to 0.
- TLBP Set to 1 if the exception is a TLB parity error; otherwise set to 0.
- ICP Set to 1 if the exception is an instruction cache parity error; otherwise set to 0.
- DCSP Set to 1 if the exception is a data cache parity error that resulted during a DCU Search operation; otherwise set to 0. See *Data Cache Parity Operations* on page 125.
- DCFP Set to 1 if the exception is a data cache parity error that resulted during a DCU Flush operation; otherwise set to 0. See *Data Cache Parity Operations* on page 125.
- IMPE Set to 1 if MCSR[MCS] is set (or would be, if it were not already set) and MSR[ME] = 0; otherwise set to 0. When set, this bit indicates that a Machine Check exception happened while Machine Check interrupts were disabled.

See *Machine Check Interrupts* on page 155 for more information on the handling of Machine Check interrupts within the PPC440x5 core.

**Programming Note:** If a Instruction Synchronous Machine Check exception occurs (i.e. an error occurs on the PLB transfer that is intended to fill a line in the instruction cache, any data associated with the exception will *not* be placed into the instruction cache. On the other hand, if a Data Asynchronous Machine Check exception occurs due to a PLB error during a cacheable read operation, the data associated with the exception

will be placed into the data cache, and could subsequently be loaded into a register. Similarly, if a Data Asynchronous Machine Check exception due to a PLB error occurs during a caching inhibited read operation, the data associated with the exception will be read into a register. Data Asynchronous Machine Check exceptions resulting from parity errors may or may not corrupt a GPR value, depending on the setting of the CCR0[PRE] field. See *Data Cache Parity Operations* on page 125.

**Programming Note:** Since a **dcbz** instruction establishes a real address in the data cache without actually reading the block of data from memory, it is possible for a delayed Data Machine Check exception to occur if and when a line established by a **dcbz** instruction is cast-out of the data cache and written to memory, if the address of the cache line is not valid within the system implementation.

### 6.5.3 Data Storage Interrupt

A Data Storage interrupt *may* occur when no higher priority exception exists and a Data Storage exception is presented to the interrupt mechanism. The PPC440x5 core includes four types of Data Storage exception. They are:

#### Read Access Control exception

A Read Access Control exception is caused by one of the following:

- While in user mode (MSR[PR] = 1), a load, **icbi**, **icbt**, **dcbst**, **dcbf**, **dcbt**, or **dcbtst** instruction attempts to access a location in storage that is not enabled for read access in user mode (that is, the TLB entry associated with the memory page being accessed has UR=0).
- While in supervisor mode (MSR[PR] = 0), a load, **icbi**, **icbt**, **dcbst**, **dcbf**, **dcbt**, or **dcbtst** instruction attempts to access a location in storage that is not enabled for read access in supervisor mode (that is, the TLB entry associated with the memory page being accessed has SR=0).

**Programming Note:** The instruction cache management instructions **icbi** and **icbt** are treated as “loads” from the addressed byte with respect to address translation and protection. These instruction cache management instructions use MSR[DS] rather than MSR[IS] to determine translation for their target effective address. Similarly, they use the read access control field (UR or SR) rather than the execute access control field (UX or SX) of the TLB entry to determine whether a Data Storage exception should occur. Instruction Storage exceptions and Instruction TLB Miss exceptions are associated with the *fetching* of instructions not with the *execution* of instructions. Data Storage exceptions and Data TLB Miss exceptions are associated with the *execution* of instruction cache management instructions, as well as with the execution of load, store, and data cache management instructions.

#### Write Access Control exception

A Write Access Control exception is caused by one of the following:

- While in user mode (MSR[PR] = 1), a store, **dcbz**, or **dcbi** instruction attempts to access a location in storage that is not enabled for write access in user mode (that is, the TLB entry associated with the memory page being accessed has UW=0).



- While in supervisor mode ( $MSR[PR] = 0$ ), a store, **dcbz**, or **dcbi** instruction attempts to access a location in storage that is not enabled for write access in supervisor mode (that is, the TLB entry associated with the memory page being accessed has  $SW=0$ ).

### Byte Ordering exception

A Byte Ordering exception will occur when a floating-point unit or auxiliary processor is attached to the PPC440x5 core, and a floating-point or auxiliary processor load or store instruction attempts to access a memory page with a byte order which is not supported by the attached processor. Whether or not a given load or store instruction type is supported for a given byte order is dependent on the implementation of the floating-point or auxiliary processor. All integer load and store instructions supported by the PPC440x5 core are supported for both big endian and little endian memory pages.

### Cache Locking exception

A Cache Locking exception is caused by one of the following:

- While in user mode ( $MSR[PR] = 1$ ) with  $MMUCR[IULXE]=1$ , execution of an **icbi** instruction is attempted. The exception occurs whether or not the cache line targeted by the **icbi** instruction is actually locked in the instruction cache.
- While in user mode ( $MSR[PR] = 1$ ) with  $MMUCR[DULXE]=1$ , execution of a **dcbf** instruction is attempted. The exception occurs whether or not the cache line targeted by the **dcbf** instruction is actually locked in the data cache.

See *Instruction and Data Caches* on page 93 and *Memory Management Unit Control Register (MMUCR)* on page 143 for more information on cache locking and Cache Locking exceptions, respectively.

If a **stwcx**. instruction causes a Write Access Control exception, but the processor does not have the reservation from a **lwarx** instruction, then a Data Storage interrupt does not occur and the instruction completes, updating  $CR[CR0]$  to indicate the failure of the store due to the lost reservation.

**If a Data Storage exception occurs on any of the following instructions, then the instruction is treated as a no-op, and a Data Storage interrupt does not occur.**

- **lswx** or **stswx** with a length of zero (although the target register of **lswx** will still be undefined, as it is whether or not a Data Storage exception occurs)
- **icbt**
- **dcbt**
- **dcbtst**

For all other instructions, if a Data Storage exception occurs, then execution of the instruction causing the exception is suppressed, a Data Storage interrupt is generated, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR2[IVO] \parallel 0b0000$ .

### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Data Storage interrupt.

### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.



**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**Data Exception Address Register (DEAR)**

If the instruction causing the Data Storage exception does so with respect to the memory page targeted by the initial effective address calculated by the instruction, then the DEAR is set to this calculated effective address. On the other hand, if the Data Storage exception only occurs due to the instruction causing the exception crossing a memory page boundary, in that the exception is with respect to the attributes of the page accessed after crossing the boundary, then the DEAR is set to the address of the first byte within that page.

For example, consider a misaligned load word instruction that targets effective address 0x00000FFF, and that the page containing that address is a 4KB page. The load word will thus cross the page boundary, and access the next page starting at address 0x00001000. If a Read Access Control exception exists within the first page (because the Read Access Control field for that page is 0), the DEAR will be set to 0x00000FFF. On the other hand, if the Read Access Control field of the first page is 1, but the same field is 0 for the next page, then the Read Access Control exception exists only for the second page and the DEAR will be set to 0x00001000. Furthermore, the load word instruction in this latter scenario will have been partially executed (see *Partially Executed Instructions* on page 158).

**Exception Syndrome Register (ESR)**

FP Set to 1 if the instruction causing the interrupt is a floating-point load or store; otherwise set to 0.

ST Set to 1 if the instruction causing the interrupt is a store, **dcbz**, or **dcbi** instruction; otherwise set to 0.

DLK<sub>0,1</sub> Set to 0b10 if an **icbi** instruction caused a Cache Locking exception; set to 0b01 if a **dcbf** instruction caused a Cache Locking exception; otherwise set to 0b00. Note that a Read Access Control exception may occur in combination with a Cache Locking exception, in which case software would need to examine the TLB entry associated with the address reported in the DEAR to determine whether both exceptions had occurred, or just a Cache Locking exception.

AP Set to 1 if the instruction causing the interrupt is an auxiliary processor load or store; otherwise set to 0.

BO Set to 1 if the instruction caused a Byte Ordering exception; otherwise set to 0. Note that a Read or Write Access Control exception may occur in combination with a Byte Ordering exception, in which case software would need to examine the TLB entry associated with the address reported in the DEAR to determine whether both exceptions had occurred, or just a Byte Ordering exception.

MCI Unchanged.

All other defined ESR bits are set to 0.

### 6.5.4 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. Note that although an Instruction Storage exception may occur during an attempt to *fetch* an instruction, such an exception is not actually presented to the interrupt mechanism until an attempt is made to *execute* that instruction. The PPC440x5 core includes one type of Instruction Storage exception. That is:

#### Execute Access Control exception

An Execute Access Control exception is caused by one of the following:

- While in user mode ( $MSR[PR] = 1$ ), an instruction fetch attempts to access a location in storage that is not enabled for execute access in user mode (that is, the TLB entry associated with the memory page being accessed has  $UX = 0$ ).
- While in supervisor mode ( $MSR[PR] = 0$ ), an instruction fetch attempts to access a location in storage that is not enabled for execute access in supervisor mode (that is, the TLB entry associated with the memory page being accessed has  $SX = 0$ ).

**Architecture Note:** The PowerPC Book-E architecture defines an additional Instruction Storage exception -- the Byte Ordering exception. This exception is defined to assist implementations that cannot support dynamically switching byte ordering between consecutive instruction fetches and/or cannot support a given byte order at all. The PPC440x5 core however supports instruction fetching from both big endian and little endian memory pages, so this exception cannot occur.

When an Instruction Storage interrupt occurs, the processor suppresses the execution of the instruction causing the Instruction Storage exception, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] || IVOR3[IVO] || 0b0000$ .

#### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Instruction Storage interrupt.

#### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

#### Exception Syndrome Register (ESR)

BO Set to 0.

MCI Unchanged.

All other defined ESR bits are set to 0.

### 6.5.5 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. An External Input exception is caused by the activation of an asynchronous input to the PPC440x5 core. Although the only mask for this interrupt type within the core is the MSR[EE] bit, system implementations typically provide an alternative means for independently masking the interrupt requests from the various devices which collectively may activate the core's External Input interrupt request input.

**Note:** MSR[EE] also enables the External Input and Fixed-Interval Timer interrupts.

When an External Input interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR4[IVO] || 0b0000.

#### Save/Restore Register 0 (SRR0)

Set to the effective address of the next instruction to be executed.

#### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**Programming Note:** Software is responsible for taking any action(s) that are required by the implementation in order to clear any External Input exception status (such that the External Input interrupt request input signal is deasserted) before reenabling MSR[EE], in order to avoid another, redundant External Input interrupt.

### 6.5.6 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 occurs if execution of any of the following is attempted:

- An integer load or store instruction that references a data storage operand that is not aligned on an operand-sized boundary, when CCR0[FLSTA] is 1. Load and store multiple instructions are considered to reference word operands, and hence word-alignment is required for the target address of these instructions when CCR0[FLSTA] is 1. Load and store string instructions are considered to reference byte operands, and hence they cannot cause an Alignment exception due to CCR0[FLSTA] being 1, regardless of the target address alignment.
- A floating-point or auxiliary processor load or store instruction that references a data storage operand that crosses a quadword (16 byte) boundary.
- A floating-point or auxiliary processor load or store instruction that references a data storage operand that is not aligned on an operand-sized boundary, when the attached processing unit indicates to the PPC440x5 core that the instruction requires operand-alignment.

- A floating-point or auxiliary processor load or store instruction that references a data storage operand that is not aligned on a word boundary, when the attached processing unit indicates to the PPC440x5 core that the instruction requires word-alignment.
- A **dcbz** instruction that targets a memory page that is either write-through required or caching inhibited.

If a **stwcx.** instruction causes an Alignment exception, and the processor does not have the reservation from a **lwarx** instruction, then an Alignment interrupt still occurs.

**Programming Note:** The architecture does not support the use of an unaligned effective address by the **lwarx** and **stwcx.** instructions. If an Alignment interrupt occurs due to the attempted execution of one of these instructions, the Alignment interrupt handler must not attempt to emulate the instruction, but instead should treat the instruction as a programming error.

When an Alignment interrupt occurs, the processor suppresses the execution of the instruction causing the Alignment exception, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR5[IVO] \parallel 0b0000$ .

#### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Alignment interrupt.

#### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

#### Data Exception Address Register (DEAR)

Set to the effective address of the target data operand as calculated by the instruction causing the Alignment exception. Note that for **dcbz**, this effective address is not necessarily the address of the first byte of the targeted cache block, but could be the address of any byte within the block (it will be the address calculated by the **dcbz** instruction).

#### Exception Syndrome Register (ESR)

FP Set to 1 if the instruction causing the interrupt is a floating-point load or store; otherwise set to 0.

ST Set to 1 if the instruction causing the interrupt is a store, **dcbz**, or **dcbi** instruction; otherwise set to 0.

AP Set to 1 if the instruction causing the interrupt is an auxiliary processor load or store; otherwise set to 0.

All other defined ESR bits are set to 0.

### 6.5.7 Program Interrupt

A Program interrupt occurs when no higher priority exception exists, a Program exception is presented to the interrupt mechanism, and -- for the Floating-Point Enabled form of Program exception only -- MSR[FE0,FE1] is non-zero. The PPC440x5 core includes following types of Program exception:

### Illegal Instruction exception

An Illegal Instruction exception occurs when execution is attempted of any of the following kinds of instructions:

- a reserved-illegal instruction
- when  $MSR[PR] = 1$  (user mode), an **mtspr** or **mfspr** that specifies an SPRN value with  $SPRN_5 = 0$  (user-mode accessible) that represents an unimplemented Special Purpose Register. For **mtspr**, this includes any SPR number other than the XER, LR, CTR, or USPRG0. For **mfspr**, this includes any SPR number other than the ones listed for **mtspr**, plus SPRG4-7, TBH, and TBL.
- a defined instruction which is not implemented within the PPC440x5 core, and which is not a floating-point instruction. This includes all instructions that are defined for 64-bit implementations only, as well as **tlbiva** and **mfapidi** (see the PowerPC Book-E specification)
- a defined floating-point instruction that is not recognized by an attached floating-point unit (or when no such floating-point unit is attached)
- an allocated instruction that is not implemented within the PPC440x5 core and which is not recognized by an attached auxiliary processor (or when no such auxiliary processor is attached)

See *Instruction Classes* on page 52 for more information on the PPC440x5 core's support for defined and allocated instructions.

### Privileged Instruction exception

A Privileged Instruction exception occurs when  $MSR[PR] = 1$  and execution is attempted of any of the following kinds of instructions:

- a privileged instruction
- an **mtspr** or **mfspr** instruction that specifies an SPRN value with  $SPRN_5 = 1$  (a Privileged Instruction exception occurs regardless of whether or not the SPR referenced by the SPRN value is defined)

### Trap exception

A Trap exception occurs when any of the conditions specified in a **tw** or **twi** instruction are met. However, if Trap debug events are enabled ( $DBCR0[TRAP]=1$ ), internal debug mode is enabled ( $DBCR0[IDM]=1$ ), and Debug interrupts are enabled ( $MSR[DE]=1$ ), then a Trap exception will cause a Debug interrupt to occur, rather than a Program interrupt.

See *Debug Facilities* on page 213 for more information on Trap debug events.

### Unimplemented Operation exception

An Unimplemented Operation exception occurs when execution is attempted of any of the following kinds of instructions:

- a defined floating-point instruction that is recognized but not supported by an attached floating-point unit, when floating-point instruction processing is enabled ( $MSR[FP]=1$ ).
- an allocated instruction that is not implemented within the PPC440x5 core, and is recognized but not supported by an attached auxiliary processor, when auxiliary processor instruction processing is enabled. The enabling of auxiliary processor instruction processing is implementation-dependent.

### Floating-Point Enabled exception

A Floating-Point Enabled exception occurs when the execution or attempted execution of a defined floating-point instruction causes  $FPSCR[FEX]$  to be set to 1, in an attached floating-point unit.  $FPSCR[FEX]$  is the Floating-Point Status and Control Register Floating-Point Enabled Exception Summary bit (see the user's manual for the floating-point unit implementation for more

details).

If MSR[FE0,FE1] is non-zero when the Floating-Point Enabled exception is presented to the interrupt mechanism, then a Program interrupt will occur, and the interrupt processing registers will be updated as described below. If MSR[FE0,FE1] are both 0, however, then a Program interrupt will *not* occur and the instruction associated with the exception will execute according to the definition of the floating-point unit (see the user's manual for the floating-point unit implementation). If and when MSR[FE0,FE1] are subsequently set to a non-zero value, and the Floating-Point Enabled exception is still being presented to the interrupt mechanism (that is, FPSCR[FEX] is still set), then a "delayed" Program interrupt will occur, updating the interrupt processing registers as described below.

See *Synchronous, Imprecise Interrupts* on page 154 for more information on this special form of "delayed" Floating-Point Enabled exception.

### Auxiliary Processor Enabled exception

An Auxiliary Processor Enabled exception may occur due to the execution or attempted execution of an allocated instruction that is not implemented within the PPC440x5 core, but is recognized and supported by an attached auxiliary processor. The cause of such an exception is implementation-dependent. See the user's manual for the auxiliary processor implementation for more details.

When a Program interrupt occurs, the processor suppresses the execution of the instruction causing the Program exception (for all cases except the "delayed" form of Floating-Point Enabled exception described above), the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR6[IVO] || 0b0000.

### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Program interrupt, for all cases except the "delayed" form of Floating-Point Enabled exception described above.

For the special case of the delayed Floating-Point Enabled exception, where the exception was already being presented to the interrupt mechanism at the time MSR[FE0,FE1] was changed from 0 to a non-zero value, SRR0 is set to the address of the instruction that would have executed after the MSR-changing instruction. If the instruction which set MSR[FE0,FE1] was **rfi**, **rftci**, or **rfmci**, then CSRR0 is set to the address to which the **rftci**, **rftci**, or **rfmci** was returning, and not to the address of the instruction which was sequentially after the **rftci**, **rftci**, or **rfmci**.

### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**Exception Syndrome Register (ESR)**

PIL	Set to 1 for an Illegal Instruction exception; otherwise set to 0
PPR	Set to 1 for a Privileged Instruction exception; otherwise set to 0
PTR	Set to 1 for a Trap exception; otherwise set to 0
PUO	Set to 1 for an Unimplemented Operation exception; otherwise set to 0
FP	Set to 1 if the instruction causing the interrupt is a floating-point instruction; otherwise set to 0.
AP	Set to 1 if the instruction causing the interrupt is an auxiliary processor instruction; otherwise set to 0.
PIE	Set to 1 if a “delayed” form of Floating-point Enabled exception type Program interrupt; otherwise set to 0. The setting of ESR[PIE] to 1 indicates to the Program interrupt handler that the interrupt was imprecise due to being caused by the changing of MSR[FE0,FE1] and not directly by the execution of the floating-point instruction which caused the exception by setting FPSCR[FEX]. Thus the Program interrupt handler can recognize that SRR0 contains the address of the instruction after the MSR-changing instruction, and not the address of the instruction that caused the Floating-Point Enabled exception.
PCRE	Set to 1 if a Floating-Point Enabled exception and the floating-point instruction which caused the exception was a CR-updating instruction. Note that ESR[PCRE] is undefined if ESR[PIE] is 1.
PCMP	Set to 1 if a Floating-Point Enabled exception and the instruction which caused the exception was a floating-point compare instruction. Note that ESR[PCMP] is undefined if ESR[PIE] is 1.
PCRF	Set to the number of the CR field (0 - 7) which was to be updated, if a Floating-Point Enabled exception and the floating-point instruction which caused the exception was a CR-updating instruction. Note that ESR[PCRF] is undefined if ESR[PIE] is 1.

All other defined ESR bits are set to 0.

**Programming Note:** The ESR[PCRE,PCMP,PCRF] fields are provided to assist the Program interrupt handler with the emulation of part of the function of the various floating-point CR-updating instructions, when any of these instructions cause a precise (“non-delayed”) Floating-Point Enabled exception type Program interrupt. The PowerPC Book-E floating-point architecture defines that when such exceptions occur, the CR is to be updated even though the rest of the instruction execution may be suppressed. The PPC440x5 core, however, does not support such CR updates when the instruction which is supposed to cause the update is being suppressed due to the occurrence of a synchronous, precise interrupt. Instead, the PPC440x5 core records in the ESR[PCRE,PCMP,PCRF] fields information about the instruction causing the interrupt, to assist the Program interrupt handler software in performing the appropriate CR update manually.

**6.5.8 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 which is recognized by an attached floating-point unit, and MSR[FP]=0.



When a Floating-Point Unavailable interrupt occurs, the processor suppresses the execution of the instruction causing the Floating-Point Unavailable exception, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR7[IVO] \parallel 0b0000$ .

**Save/Restore Register 0 (SRR0)**

Set to the effective address of the next instruction causing the Floating-Point Unavailable interrupt.

**Save/Restore Register 1 (SRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**6.5.9 System Call Interrupt**

A System Call interrupt occurs when no higher priority exception exists and a system call (**sc**) instruction is executed.

When a System Call interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR8[IVO] \parallel 0b0000$ .

**Save/Restore Register 0 (SRR0)**

Set to the effective address of the instruction after the system call instruction.

**Save/Restore Register 1 (SRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**6.5.10 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 which is not implemented within the PPC440x5 core but which is recognized by an attached auxiliary processor, and auxiliary processor instruction processing is not enabled. The enabling of auxiliary processor instruction processing is implementation-dependent. See the user's manual for the attached auxiliary processor.

When an Auxiliary Processor Unavailable interrupt occurs, the processor suppresses the execution of the instruction causing the Auxiliary Processor Unavailable exception, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR9[IVO] \parallel 0b0000$ .



**Save/Restore Register 0 (SRR0)**

Set to the effective address of the next instruction causing the Auxiliary Processor Unavailable interrupt.

**Save/Restore Register 1 (SRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**6.5.11 Decrementer Interrupt**

A Decrementer interrupt occurs when no higher priority exception exists, a Decrementer exception exists (TSR[DIS] = 1), and the interrupt is enabled (TCR[DIE] = 1 and MSR[EE] = 1). See *Timer Facilities* on page 203 for more information on Decrementer exceptions.

**Note:** MSR[EE] also enables the External Input and Fixed-Interval Timer interrupts.

When a Decrementer interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR10[IVO] || 0b0000.

**Save/Restore Register 0 (SRR0)**

Set to the effective address of the next instruction to be executed.

**Save/Restore Register 1 (SRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**Programming Note:** Software is responsible for clearing the Decrementer exception status by writing to TSR[DIS], prior to reenabling MSR[EE], in order to avoid another, redundant Decrementer interrupt.

**6.5.12 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). See *Timer Facilities* on page 203 for more information on Fixed Interval Timer exceptions.

**Note:** MSR[EE] also enables the External Input and Decrementer interrupts.

When a Fixed interval Timer interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR11[IVO] || 0b0000.

**Save/Restore Register 0 (SRR0)**

Set to the effective address of the next instruction to be executed.

**Save/Restore Register 1 (SRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

CE, ME, DE Unchanged.

All other MSR bits set to 0.

**Programming Note:** Software is responsible for clearing the Fixed Interval Timer exception status by writing to TSR[FIS], prior to reenabling MSR[EE], in order to avoid another, redundant Fixed Interval Timer interrupt.

**6.5.13 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). See *Timer Facilities* on page 203 for more information on Watchdog Timer exceptions.

**Note:** MSR[CE] also enables the Critical Input interrupt.

When a Watchdog Timer interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address IVPR[IVP] || IVOR12[IVO] || 0b0000.

**Critical Save/Restore Register 0 (CSRR0)**

Set to the effective address of the next instruction to be executed.

**Critical Save/Restore Register 1 (CSRR1)**

Set to the contents of the MSR at the time of the interrupt.

**Machine State Register (MSR)**

ME Unchanged.

All other MSR bits set to 0.

**Programming Note:** Software is responsible for clearing the Watchdog Timer exception status by writing to TSR[WIS], prior to reenabling MSR[CE], in order to avoid another, redundant Watchdog Timer interrupt.

**6.5.14 Data TLB Error Interrupt**

A Data TLB Error interrupt *may* occur when no higher priority exception exists and a Data TLB Miss exception is presented to the interrupt mechanism. A Data TLB Miss exception occurs when a load, store, **icbi**, **icbt**, **dcbst**, **dcbf**, **dcbz**, **dcbi**, **dcbt**, or **dcbtst** instruction attempts to access a virtual address for which a valid TLB entry does not exist. See *Memory Management* on page 129 for more information on the TLB.

**Programming Note:** The instruction cache management instructions **icbi** and **icbt** are treated as “loads” from the addressed byte with respect to address translation and protection, and therefore use MSR[DS] rather than MSR[IS] as part

of the calculated virtual address when searching the TLB to determine translation for their target storage address. Instruction TLB Miss exceptions are associated with the *fetching* of instructions not with the *execution* of instructions. Data TLB Miss exceptions are associated with the *execution* of instruction cache management instructions, as well as with the execution of load, store, and data cache management instructions.

If a **stwcx.** instruction causes a Data TLB Miss exception, and the processor does not have the reservation from a **lwarx** instruction, then a Data TLB Error interrupt still occurs.

**If a Data TLB Miss exception occurs on any of the following instructions, then the instruction is treated as a no-op, and a Data TLB Error interrupt does not occur.**

- **lswx** or **stswx** with a length of zero (although the target register of **lswx** will be undefined)
- **icbt**
- **dcbt**
- **dcbtst**

For all other instructions, if a Data TLB Miss exception occurs, then execution of the instruction causing the exception is suppressed, a Data TLB Error interrupt is generated, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address `IVPR[IVP] || IVOR13[IVO] || 0b0000`.

#### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Data TLB Error interrupt.

#### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

#### Data Exception Address Register (DEAR)

If the instruction causing the Data TLB Miss exception does so with respect to the memory page targeted by the initial effective address calculated by the instruction, then the DEAR is set to this calculated effective address. On the other hand, if the Data TLB Miss exception only occurs due to the instruction causing the exception crossing a memory page boundary, in that the missing TLB entry is for the page accessed after crossing the boundary, then the DEAR is set to the address of the first byte within that page.

As an example, consider a misaligned load word instruction that targets effective address `0x00000FFF`, and that the page containing that address is a 4KB page. The load word will thus cross the page boundary, and attempt to access the next page starting at address `0x00001000`. If a valid TLB entry does not exist for the first page, then the DEAR will be set to `0x00000FFF`. On the other hand, if a valid TLB entry *does* exist for the first page, but not for the second, then the DEAR will be set to `0x00001000`. Furthermore, the load word instruction in this latter scenario will have been partially executed (see *Partially Executed Instructions* on page 158).

### Exception Syndrome Register (ESR)

FP	Set to 1 if the instruction causing the interrupt is a floating-point load or store; otherwise set to 0.
ST	Set to 1 if the instruction causing the interrupt is a store, <b>dcbz</b> , or <b>dcbi</b> instruction; otherwise set to 0.
AP	Set to 1 if the instruction causing the interrupt is an auxiliary processor load or store; otherwise set to 0.
MCI	Unchanged.

All other defined ESR bits are set to 0.

#### 6.5.15 Instruction TLB Error Interrupt

An Instruction TLB Error interrupt occurs when no higher priority exception exists and an Instruction TLB Miss exception is presented to the interrupt mechanism. Note that although an Instruction TLB Miss exception may occur during an attempt to *fetch* an instruction, such an exception is not actually presented to the interrupt mechanism until an attempt is made to *execute* that instruction. An Instruction TLB Miss exception occurs when an instruction fetch attempts to access a virtual address for which a valid TLB entry does not exist. See *Memory Management* on page 129 for more information on the TLB.

When an Instruction TLB Error interrupt occurs, the processor suppresses the execution of the instruction causing the Instruction TLB Miss exception, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged), and instruction execution resumes at address  $IVPR[IVP] \parallel IVOR14[IVO] \parallel 0b0000$ .

#### Save/Restore Register 0 (SRR0)

Set to the effective address of the instruction causing the Instruction TLB Error interrupt.

#### Save/Restore Register 1 (SRR1)

Set to the contents of the MSR at the time of the interrupt.

#### Machine State Register (MSR)

CE, ME, DE Unchanged.

All other MSR bits set to 0.

#### 6.5.16 Debug Interrupt

A Debug interrupt occurs when no higher priority exception exists, a Debug exception exists in the Debug Status Register (DBSR), the processor is in internal debug mode ( $DBCR0[IDM]=1$ ), and Debug interrupts are enabled ( $MSR[DE] = 1$ ). A Debug exception occurs when a debug event causes a corresponding bit in the DBSR to be set.

There are several types of Debug exception, as follows:

#### Instruction Address Compare (IAC) exception

An IAC Debug exception occurs when execution is attempted of an instruction whose address matches the IAC conditions specified by the various debug facility registers. This exception can occur regardless of debug mode, and regardless of the value of  $MSR[DE]$ .

### Data Address Compare (DAC) exception

A DAC Debug exception occurs when the DVC mechanism is not enabled, and execution is attempted of a load, store, **icbi**, **icbt**, **dcbst**, **dcbf**, **dcbz**, **dcbi**, **dcbt**, or **dcbtst** instruction whose target storage operand address matches the DAC conditions specified by the various debug facility registers. This exception can occur regardless of debug mode, and regardless of the value of MSR[DE].

**Programming Note:** The instruction cache management instructions **icbi** and **icbt** are treated as “loads” from the addressed byte with respect to Debug exceptions. IAC Debug exceptions are associated with the *fetching* of instructions not with the *execution* of instructions. DAC Debug exceptions are associated with the *execution* of instruction cache management instructions, as well as with the execution of load, store, and data cache management instructions.

### Data Value Compare (DVC) exception

A DVC Debug exception occurs when execution is attempted of a load, store, or **dcbz** instruction whose target storage operand address matches the DAC and DVC conditions specified by the various debug facility registers. This exception can occur regardless of debug mode, and regardless of the value of MSR[DE].

### Branch Taken (BRT) exception

A BRT Debug exception occurs when BRT debug events are enabled (DBCR0[BRT]=1) and execution is attempted of a branch instruction for which the branch conditions are met. This exception cannot occur in internal debug mode when MSR[DE]=0, unless external debug mode or debug wait mode is also enabled.

### Trap (TRAP) exception

A TRAP Debug exception occurs when TRAP debug events are enabled (DBCR0[TRAP]=1) and execution is attempted of a **tw** or **twi** instruction that matches any of the specified trap conditions. This exception can occur regardless of debug mode, and regardless of the value of MSR[DE].

### Return (RET) exception

A RET Debug exception occurs when RET debug events are enabled (DBCR0[RET]=1) and execution is attempted of an **rfi**, **rfdi**, or **rfmci** instruction. For **rfdi**, the RET Debug exception can occur regardless of debug mode and regardless of the value of MSR[DE]. For **rfdi** or **rfmci**, the RET Debug exception cannot occur in internal debug mode when MSR[DE]=0, unless external debug mode or debug wait mode is also enabled.

### Instruction Complete (ICMP) exception

An ICMP Debug exception occurs when ICMP debug events are enabled (DBCR0[ICMP]=1) and execution of any instruction is completed. This exception cannot occur in internal debug mode when MSR[DE]=0, unless external debug mode or debug wait mode is also enabled.

### Interrupt (IRPT) exception

An IRPT Debug exception occurs when IRPT debug events are enabled (DBCR0[IRPT]=1) and an interrupt occurs. For non-critical class interrupt types, the IRPT Debug exception can occur regardless of debug mode and regardless of the value of MSR[DE]. For critical class interrupt types, the IRPT Debug exception cannot occur in internal debug mode (regardless of the value of

MSR[DE]), unless external debug mode or debug wait mode is also enabled.

### Unconditional Debug Event (UDE) exception

A UDE Debug exception occurs when an Unconditional Debug Event is signaled over the JTAG interface to the PPC440x5 core. This exception can occur regardless of debug mode, and regardless of the value of MSR[DE].

There are four debug modes supported by the PPC440x5 core. They are: internal debug mode, external debug mode, debug wait mode, and trace mode. Debug exceptions and interrupts are affected by the debug mode(s) which are enabled at the time of the Debug exception. Debug interrupts occur only when internal debug mode is enabled, although it is possible for external debug mode and/or debug wait mode to be enabled as well. The remainder of this section assumes that internal debug mode is enabled and that external debug mode and debug wait mode are not enabled, at the time of a Debug exception.

See *Debug Facilities* on page 213 for more information on the different debug modes and the behavior of each of the Debug exception types when operating in each of the modes.

**Programming Note:** It is a programming error for software to enable internal debug mode (by setting DBCR0[IDM] to 1) while Debug exceptions are already present in the DBSR. Software must first clear all DBSR Debug exception status (that is, all fields except IDE, MRR, IAC12ATS, and IAC34ATS) before setting DBCR0[IDM] to 1.

If a **stwcx.** instruction causes a DAC or DVC Debug exception, but the processor does not have the reservation from a **lwarx** instruction, then the Debug exception is not recorded in the DBSR and a Debug interrupt does not occur. Instead, the instruction completes and updates CR[CR0] to indicate the failure of the store due to the lost reservation.

**If a DAC exception occurs on an lswx or stswx with a length of zero, then the instruction is treated as a no-op, the Debug exception is not recorded in the DBSR, and a Debug interrupt does not occur.**

**If a DAC exception occurs on an icbt, dcbt, or dcibtst instruction which is being no-op'ed due to some other reason (either the referenced cache block is in a caching inhibited memory page, or a Data Storage or Data TLB Miss exception occurs), then the Debug exception is not recorded in the DBSR and a Debug interrupt does not occur. On the other hand, if the icbt, dcbt, or dcibtst instruction is not being no-op'ed for one of these other reasons, the DAC Debug exception does occur and is handled in the same fashion as other DAC Debug exceptions (see below).**

For all other cases, when a Debug exception occurs, it is immediately presented to the interrupt handling mechanism. A Debug interrupt will occur immediately if MSR[DE] is 1, and the interrupt processing registers will be updated as described below. If MSR[DE] is 0, however, then the exception condition remains set in the DBSR. If and when MSR[DE] is subsequently set to 1, and the exception condition is still present in the DBSR, a "delayed" Debug interrupt will then occur either as a synchronous, imprecise interrupt, or as an asynchronous interrupt, depending on the type of Debug exception.

When a Debug interrupt occurs, the interrupt processing registers are updated as indicated below (all registers not listed are unchanged) and instruction execution resumes at address IVPR[IVP] || IVOR15[IVO] || 0b0000.

### Critical Save/Restore Register 0 (CSRR0)

For Debug exceptions that occur while Debug interrupts are enabled (MSR[DE] = 1), CSRR0 is set as follows:

- For IAC, BRT, TRAP, and RET Debug exceptions, set to the address of the instruction causing the Debug interrupt. Execution of the instruction causing the Debug exception is suppressed, and the interrupt is synchronous and precise.
- For DAC and DVC Debug exceptions, if DBCR2[**DAC12A**] is 0, set to the address of the instruction causing the Debug interrupt. Execution of the instruction causing the Debug exception is suppressed, and the interrupt is synchronous and precise.

If DBCR2[**DAC12A**] is 1, however, then DAC and DVC Debug exceptions are handled asynchronously, and CSRR0 is set to the address of the instruction that would have executed next had the Debug interrupt not occurred. This could either be the address of the instruction causing the DAC or DVC Debug exception, or the address of a subsequent instruction.

- For ICMP Debug exceptions, set to the address of the next instruction to be executed (the instruction after the one whose completion caused the ICMP Debug exception). The interrupt is synchronous and precise.

Since the ICMP Debug exception does not suppress the execution of the instruction causing the exception, but rather allows it to complete before causing the interrupt, the behavior of the interrupt is different in the special case where the instruction causing the ICMP Debug exception is itself setting MSR[**DE**] to 0. In this case, the interrupt will be delayed and will occur if and when MSR[**DE**] is again set to 1, assuming DBSR[**ICMP**] is still set. If the Debug interrupt occurs in this fashion, it will be synchronous and imprecise, and CSRR0 will be set to the address of the instruction after the one which set MSR[**DE**] to 1 (not the one which originally caused the ICMP Debug exception and in so doing set MSR[**DE**] to 0). If the instruction which set MSR[**DE**] to 1 was **rfi**, **rfci**, or **rfmci**, then CSRR0 is set to the address to which the **rfi**, **rfci**, or **rfmci** was returning, and not to the address of the instruction which was sequentially after the **rfi**, **rfci**, or **rfmci**.

- For IRPT Debug exceptions, set to the address of the first instruction in the interrupt handler associated with the interrupt type that caused the IRPT Debug exception. The interrupt is asynchronous.
- For UDE Debug exceptions, set to the address of the instruction that would have executed next if the Debug interrupt had not occurred. The interrupt is asynchronous.

For all Debug exceptions that occur while Debug interrupts are disabled (MSR[**DE**] = 0), the Debug interrupt will be delayed and will occur if and when MSR[**DE**] is again set to 1, assuming the Debug exception status is still set in the DBSR. If the Debug interrupt occurs in this fashion, CSRR0 is set to the address of the instruction after the one which set MSR[**DE**]. If the instruction which set MSR[**DE**] was **rfi**, **rfci**, or **rfmci**, then CSRR0 is set to the address to which the **rfi**, **rfci**, or **rfmci** was returning, and not to the address of the instruction which was sequentially after the **rfi**, **rfci**, or **rfmci**. The interrupt is either synchronous and imprecise, or asynchronous, depending on the type of Debug exception, as follows:

- For IAC and RET Debug exceptions, the interrupt is synchronous and imprecise.
- For BRT Debug exceptions, this scenario cannot occur. BRT Debug exceptions are not recognized when MSR[**DE**]=0 if operating in internal debug mode.
- For TRAP Debug exceptions, the Debug interrupt is synchronous and imprecise. However, under these conditions (TRAP Debug exception occurring while MSR[**DE**] is 0), the attempted execution of the trap instruction for which one or more of the trap conditions is met will itself lead to a Trap exception type Program interrupt. The corresponding Debug interrupt which will occur later if and when Debug interrupts are enabled will be *in addition* to the Program interrupt.
- For DAC and DVC Debug exceptions, if DBCR2[**DAC12A**] is 0, then the interrupt is synchronous and imprecise. If DBCR2[**DAC12A**] is 1, then the interrupt is asynchronous.



- For ICMP Debug exceptions, this scenario cannot occur in this fashion. ICMP Debug exceptions are not recognized when MSR[DE]=0 if operating in internal debug mode. However, a similar scenario can occur when MSR[DE] is 1 at the time of the ICMP Debug exception, but the instruction whose completion is causing the exception is itself setting MSR[DE] to 0. This scenario is described above in the subsection on the ICMP Debug exception for which MSR[DE] is 1 at the time of the exception. In that scenario, the interrupt is synchronous and imprecise.
- For IRPT and UDE Debug exceptions, the interrupt is asynchronous.

### Critical Save/Restore Register 1 (CSRR1)

Set to the contents of the MSR at the time of the interrupt.

### Machine State Register (MSR)

ME Unchanged.

All other MSR bits set to 0.

## 6.6 Interrupt Ordering and Masking

It is possible for multiple exceptions to exist simultaneously, each of which could cause the generation of an interrupt. Furthermore, the PowerPC Book-E architecture does not provide for the generation of more than one interrupt of the same class (critical or non-critical) at a time. Therefore, the architecture defines that interrupts are ordered with respect to each other, and provides a masking mechanism for 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). However, 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 will then finally be generated.

All asynchronous interrupt types can be masked. Machine Check interrupts can be masked, as well. In addition, certain synchronous interrupt types can be masked. The two synchronous interrupt types which can be masked are the Floating-Point Enabled exception type Program interrupt (masked by MSR[FE0,FE1]), and the IAC, DAC, DVC, RET, and ICMP exception type Debug interrupts (masked by MSR[DE]).

**Architecture Note:** When an otherwise synchronous, *precise* interrupt type is “delayed” in this fashion via masking, and the interrupt type is later enabled, the interrupt that is then generated due to the exception event that occurred while the interrupt type was disabled is then considered a synchronous, *imprecise* class of interrupt.

In order to prevent a subsequent interrupt from causing the state information (saved in SRR0/SRR1, CSRR0/CSRR1, or MCSRR0/MCSRR1) from a previous interrupt to be overwritten and lost, the PPC440x5 core performs certain functions. As a first step, upon any non-critical class interrupt, the processor automatically disables any further asynchronous, non-critical class interrupts (External Input, Decrementer, and Fixed Interval Timer) by clearing MSR[EE]. Likewise, upon any critical class interrupt, hardware automatically disables any further asynchronous interrupts of either class (critical and non-critical) by clearing MSR[CE] and MSR[DE], in addition to MSR[EE]. The additional interrupt types that are disabled by the clearing of MSR[CE,DE] are the Critical Input, Watchdog Timer, and Debug interrupts. For machine check interrupts, the processor automatically disables all maskable interrupts by clearing MSR[ME] as well as MSR[EE,CE,DE].



This first step of clearing MSR[EE] (and MSR[CE,DE] for critical class interrupts, and MSR[ME] for machine checks) prevents any subsequent asynchronous interrupts from overwriting the relevant save/restore registers (SRR0/SRR1, CSRR0/CSRR1, or MCSRR0/MCSRR1), prior to software being able to save their contents. The processor also automatically clears, on any interrupt, MSR[WE,PR,FP,FE0,FE1,IS,DS]. The clearing of these bits assists in the avoidance of subsequent interrupts of certain other types. However, *guaranteeing* that these interrupt types do not occur and thus do not overwrite the save/restore registers also requires the cooperation of system software. Specifically, system software must avoid the execution of instructions that could cause (or enable) a subsequent interrupt, if the contents of the save/restore registers have not yet been saved.

### 6.6.1 Interrupt Ordering Software Requirements

The following list identifies the actions that system software must *avoid*, prior to having saved the save/restore registers' contents:

- Reenabling of MSR[EE] (or MSR[CE,DE] in critical class interrupt handlers)

This prevents any asynchronous interrupts, as well as (in the case of MSR[DE]) any Debug interrupts (which include both synchronous and asynchronous types).

- Branching (or sequential execution) to addresses not mapped by the TLB, or mapped without execute access permission

This prevents Instruction Storage and Instruction TLB Error interrupts.

- Load, store, or cache management instructions to addresses not mapped by the TLB or not having the necessary access permission (read or write)

This prevents Data Storage and Data TLB Error interrupts.

- Execution of system call (**sc**) or trap (**tw**, **twi**) instructions

This prevents System Call and Trap exception type Program interrupts.

- Execution of any floating-point instructions

This prevents Floating-Point Unavailable interrupts. Note that this interrupt would occur upon the execution of any floating-point instruction, due to the automatic clearing of MSR[FP]. However, even if software were to re-enable MSR[FP], floating-point instructions must still be avoided in order to prevent Program interrupts due to the possibility of Floating-Point Enabled and/or Unimplemented Operation exceptions.

- Reenabling of MSR[PR]

This prevents Privileged Instruction exception type Program interrupts. Alternatively, software could re-enable MSR[PR], but avoid the execution of any privileged instructions.

- Execution of any Auxiliary Processor instructions that are not implemented in the PPC440x5 core

This prevents Auxiliary Processor Unavailable interrupts, as well as Auxiliary Processor Enabled and Unimplemented Operation exception type Program interrupts. Note that the auxiliary processor instructions that are implemented within the PPC440x5 core do not cause any of these types of exceptions, and can therefore be executed prior to software having saved the save/restore registers' contents.

- Execution of any illegal instructions, or any defined instructions not implemented within the PPC440x5 core (64-bit instructions, **tlbiva**, **mfapidi**)

This prevents Illegal Instruction exception type Program interrupts.

- Execution of any instruction that could cause an Alignment interrupt

This prevents Alignment interrupts. See *Alignment Interrupt* on page 179 for a complete list of instructions that may cause Alignment interrupts.

- In the Machine Check handler, use of the caches and TLBs until any detected parity errors have been corrected.

This will avoid additional parity errors.

It is not necessary for hardware or software to avoid critical class interrupts from within non-critical class interrupt handlers (and hence the processor does not automatically clear MSR[CE,ME,DE] upon a non-critical interrupt), since the two classes of interrupts use different pairs of save/restore registers to save the instruction address and MSR. The converse, however, is not true. That is, hardware and software must cooperate in the avoidance of both critical *and* non-critical class interrupts from within critical class interrupt handlers, even though the two classes of interrupts use different save/restore register pairs. This is because the critical class interrupt may have occurred from within a non-critical class interrupt handler, prior to the non-critical class interrupt handler having saved SRR0 and SRR1. Therefore, within the critical class interrupt handler, both pairs of save/restore registers may contain data that is necessary to the system software.

Similarly, the Machine Check handler must avoid further machine checks, as well as both critical and non-critical interrupts, since the machine check handler may have been called from within a critical or non-critical interrupt handler.

### 6.6.2 Interrupt Order

The following is a prioritized listing of the various enabled interrupt types for which exceptions might exist simultaneously:

1. Synchronous (non-debug) interrupts:
  1. Data Storage
  2. Instruction Storage
  3. Alignment
  4. Program
  5. Floating-Point Unavailable
  6. System Call
  7. Auxiliary Processor Unavailable
  8. Data TLB Error
  9. Instruction TLB Error

Only one of the above types of synchronous interrupts may have an existing exception generating it at any given time. This is guaranteed by the exception priority mechanism (see *Exception Priorities* on page 195) and the requirements of the sequential execution model defined by the PowerPC Book-E architecture.

2. Machine Check
3. Debug
4. Critical Input

5. Watchdog Timer
6. External Input
7. Fixed-Interval Timer
8. Decrementer

Even though, as indicated above, the non-critical, synchronous exception types listed under item 1 are generated with higher priority than the critical interrupt types listed in items 2-5, the fact is that these non-critical interrupts will immediately be followed by the highest priority existing critical interrupt type, without executing any instructions at the non-critical interrupt handler. This is because the non-critical interrupt types do not automatically clear MSR[ME,DE,CE] and hence do not automatically disable the critical interrupt types. In all other cases, a particular interrupt type from the above list will automatically disable any subsequent interrupts of the same type, as well as all other interrupt types that are listed below it in the priority order.

## 6.7 Exception Priorities

PowerPC Book-E requires all synchronous (precise and imprecise) interrupts to be reported in program order, as implied by the sequential execution model. The one exception to this rule is the case of multiple synchronous imprecise interrupts. Upon a synchronizing event, all previously executed instructions are required to report any synchronous imprecise interrupt-generating exceptions, and the interrupt(s) will then be generated according to the general interrupt ordering rules outlined in *Interrupt Order* on page 194. For example, if a **mtmsr** instruction causes MSR[FE0,FE1,DE] to all be set, it is possible that a previous Floating-Point Enabled exception and a previous Debug exception both are still being presented (in the FPSCR and DBSR, respectively). In such a scenario, a Floating-Point Enabled exception type Program interrupt will occur first, followed immediately by a Debug interrupt.

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 will be 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 *Interrupt Order* on page 194. The exception priority mechanism also prevents certain debug exceptions from existing in combination with certain other synchronous interrupt-generating exceptions.

This section does not define the permitted setting of multiple exceptions for which the corresponding interrupt types are disabled. The generation of exceptions for which the corresponding interrupt types are disabled will have no effect on the generation of other exceptions for which the corresponding interrupt types are enabled. Conversely, if a particular exception for which the corresponding interrupt type is enabled is shown in the following sections to be of a higher priority than another exception, the occurrence of that enabled higher priority exception will prevent the setting of the other exception, independent of whether that other exception's corresponding interrupt type is enabled or disabled.

Except as specifically noted below, only one of the exception types listed for a given instruction type will be permitted to be generated at any given time, assuming the corresponding interrupt type is enabled. The priority of the exception types are listed in the following sections ranging from highest to lowest, within each instruction type.

Finally, note that Machine Check exceptions are defined by the PowerPC architecture to be neither synchronous nor asynchronous. As such, Machine Check exceptions are not considered in the remainder of this section, which is specifically addressing the priority of synchronous interrupts.

### 6.7.1 Exception Priorities for Integer Load, Store, and Cache Management Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any integer load, store, or cache management instruction. Included in this category is the former opcode for the **icbt** instruction, which is an allocated opcode still supported by the PPC440x5 core.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

Only applies to the defined 64-bit load, store, and cache management instructions, which are not recognized by the PPC440x5 core.

5. Program (Privileged Instruction)

Only applies to the **dcbi** instruction, and only occurs if MSR[PR]=1.

6. Data TLB Error (Data TLB Miss exception)
7. Data Storage (all exception types except Byte Ordering exception)
8. Alignment (Alignment exception)
9. Debug (DAC or DVC exception)
10. Debug (ICMP exception)

### 6.7.2 Exception Priorities for Floating-Point Load and Store Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any floating-point load or store instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

This exception will occur if no floating-point unit is attached to the PPC440x5 core, or if the particular floating-point load or store instruction is not recognized by the attached floating-point unit.

5. Floating-Point Unavailable (Floating-Point Unavailable exception)

This exception will occur if an attached floating-point unit recognizes the instruction, but floating-point instruction processing is disabled (MSR[FP]=0).

6. Program (Unimplemented Operation exception)

This exception will occur if an attached floating-point unit recognizes but does not support the instruction, and floating-point instruction processing is enabled (MSR[FP]=1).

7. Data TLB Error (Data TLB Miss exception)
8. Data Storage (all exception types except Cache Locking exception)
9. Alignment (Alignment exception)
10. Debug (DAC or DVC exception)
11. Debug (ICMP exception)

### 6.7.3 Exception Priorities for Allocated Load and Store Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any allocated load or store instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

This exception will occur if no auxiliary processor unit is attached to the PPC440x5 core, or if the particular allocated load or store instruction is not recognized by the attached auxiliary processor.

5. Program (Privileged Instruction exception)

This exception will occur if an attached auxiliary processor unit recognizes the instruction and indicates that the instruction is privileged, but MSR[PR]=1.

6. Auxiliary Processor Unavailable (Auxiliary Processor Unavailable exception)

This exception will occur if an attached auxiliary processor recognizes the instruction, but indicates that auxiliary processor instruction processing is disabled (whether or not auxiliary processor instruction processing is enabled is implementation-dependent).

7. Program (Unimplemented Operation exception)

This exception will occur if an attached auxiliary processor recognizes but does not support the instruction, and also indicates that auxiliary processor instruction processing is enabled (whether or not auxiliary processor instruction processing is enabled is implementation-dependent).

8. Data TLB Error (Data TLB Miss exception)
9. Data Storage (all exception types except Cache Locking exception)
10. Alignment (Alignment exception)
11. Debug (DAC or DVC exception)
12. Debug (ICMP exception)

### 6.7.4 Exception Priorities for Floating-Point Instructions (Other)

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any floating-point instruction other than a load or store.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)

3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

This exception will occur if no floating-point unit is attached to the PPC440x5 core, or if the particular floating-point instruction is not recognized by the attached floating-point unit.

5. Floating-Point Unavailable (Floating-Point Unavailable exception)

This exception will occur if an attached floating-point unit recognizes the instruction, but floating-point instruction processing is disabled (MSR[FP]=0).

6. Program (Unimplemented Operation exception)

This exception will occur if an attached floating-point unit recognizes but does not support the instruction, and floating-point instruction processing is enabled (MSR[FP]=1).

7. Program (Floating-Point Enabled exception)

This exception will occur if an attached floating-point unit recognizes and supports the instruction, floating-point instruction processing is enabled (MSR[FP]=1), and the instruction sets FPSCR[FEX] to 1.

8. Debug (ICMP exception)

### 6.7.5 Exception Priorities for Allocated Instructions (Other)

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any allocated instruction other than a load or store, and which is not one of the allocated instructions implemented within the PPC440x5 core.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

This exception will occur if no auxiliary processor unit is attached to the PPC440x5 core, or if the particular allocated instruction is not recognized by the attached auxiliary processor and is not one of the allocated instructions implemented within the PPC440x5 core.

5. Program (Privileged Instruction exception)

This exception will occur if an attached auxiliary processor unit recognizes the instruction and indicates that the instruction is privileged, but MSR[PR]=1.

6. Auxiliary Processor Unavailable (Auxiliary Processor Unavailable exception)

This exception will occur if an attached auxiliary processor recognizes the instruction, but indicates that auxiliary processor instruction processing is disabled (whether or not auxiliary processor instruction processing is enabled is implementation-dependent).

7. Program (Unimplemented Operation exception)

This exception will occur if an attached auxiliary processor recognizes but does not support the instruction, and also indicates that auxiliary processor instruction processing is enabled (whether or not auxiliary processor instruction processing is enabled is implementation-dependent).

8. Program (Auxiliary Processor Enabled exception)

This exception will occur if an attached auxiliary processor recognizes and supports the instruction, indicates that auxiliary processor instruction processing is enabled, and the instruction execution results in an Auxiliary Processor Enabled exception. Whether or not auxiliary processor instruction processing is enabled is implementation-dependent, as is whether or not a given auxiliary processor instruction results in an Auxiliary Processor Enabled exception.

9. Debug (ICMP exception)

### 6.7.6 Exception Priorities for Privileged Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of any privileged instruction other than **dcbi**, **rfi**, **rfdi**, **rfmci**, or any allocated instruction not implemented within the PPC440x5 core (all of which are covered elsewhere). This list *does* cover, however, the **dccci**, **dcread**, **iccci**, and **icread** instructions, which are privileged, allocated instructions that *are* implemented within the PPC440x5 core. This list also covers the defined 64-bit privileged instructions, the **tlbiva** instruction, and the **mfapidi** instruction, all of which are not implemented by the PPC440x5 core.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

Only applies to the defined 64-bit privileged instructions, the **tlbiva** instruction, and the **mfapidi** instruction.

5. Program (Privileged Instruction exception)

Does not apply to the defined 64-bit privileged instructions, the **tlbiva** instruction, nor the **mfapidi** instruction.

6. Debug (ICMP exception)

Does not apply to the defined 64-bit privileged instructions, the **tlbiva** instruction, nor the **mfapidi** instruction.

### 6.7.7 Exception Priorities for Trap Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of a trap (**tw**, **twi**) instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Debug (TRAP exception)
5. Program (Trap exception)
6. Debug (ICMP exception)

### 6.7.8 Exception Priorities for System Call Instruction

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of a system call (**sc**) instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. System Call (System Call exception)
5. Debug (ICMP exception)

Since the System Call exception does not suppress the execution of the **sc** instruction, but rather the exception occurs once the instruction has completed, it is possible for an **sc** instruction to cause both a System Call exception and an ICMP Debug exception at the same time. In such a case, the associated interrupts will occur in the order indicated in *Interrupt Order* on page 194.

### 6.7.9 Exception Priorities for Branch Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of a branch instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Debug (BRT exception)
5. Debug (ICMP exception)

### 6.7.10 Exception Priorities for Return From Interrupt Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of an **rfi**, **rfdi**, or **rfmci** instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Debug (RET exception)
5. Debug (ICMP exception)

### 6.7.11 Exception Priorities for Preserved Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of a preserved instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)



Applies to all preserved instructions except the **mftb** instruction, which is the only preserved class instruction implemented within the PPC440x5 core.

5. Debug (ICMP exception)

Only applies to the **mftb** instruction, which is the only preserved class instruction implemented within the PPC440x5 core.

### 6.7.12 Exception Priorities for Reserved Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of a reserved instruction.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

Applies to all reserved instruction opcodes except the reserved-nop instruction opcodes.

5. Debug (ICMP exception)

Only applies to the reserved-nop instruction opcodes.

### 6.7.13 Exception Priorities for All Other Instructions

The following list identifies the priority order of the exception types that may occur within the PPC440x5 core as the result of the attempted execution of all other instructions (that is, those not covered by one of the sections 6.7.1 through 6.7.12). This includes both defined instructions and allocated instructions implemented within the PPC440x5 core.

1. Debug (IAC exception)
2. Instruction TLB Error (Instruction TLB Miss exception)
3. Instruction Storage (Execute Access Control exception)
4. Program (Illegal Instruction exception)

Applies only to the defined 64-bit instructions, as these are not implemented within the PPC440x5 core.

5. Debug (ICMP exception)

Does not apply to the defined 64-bit instructions, as these are not implemented by the PPC440x5 core.



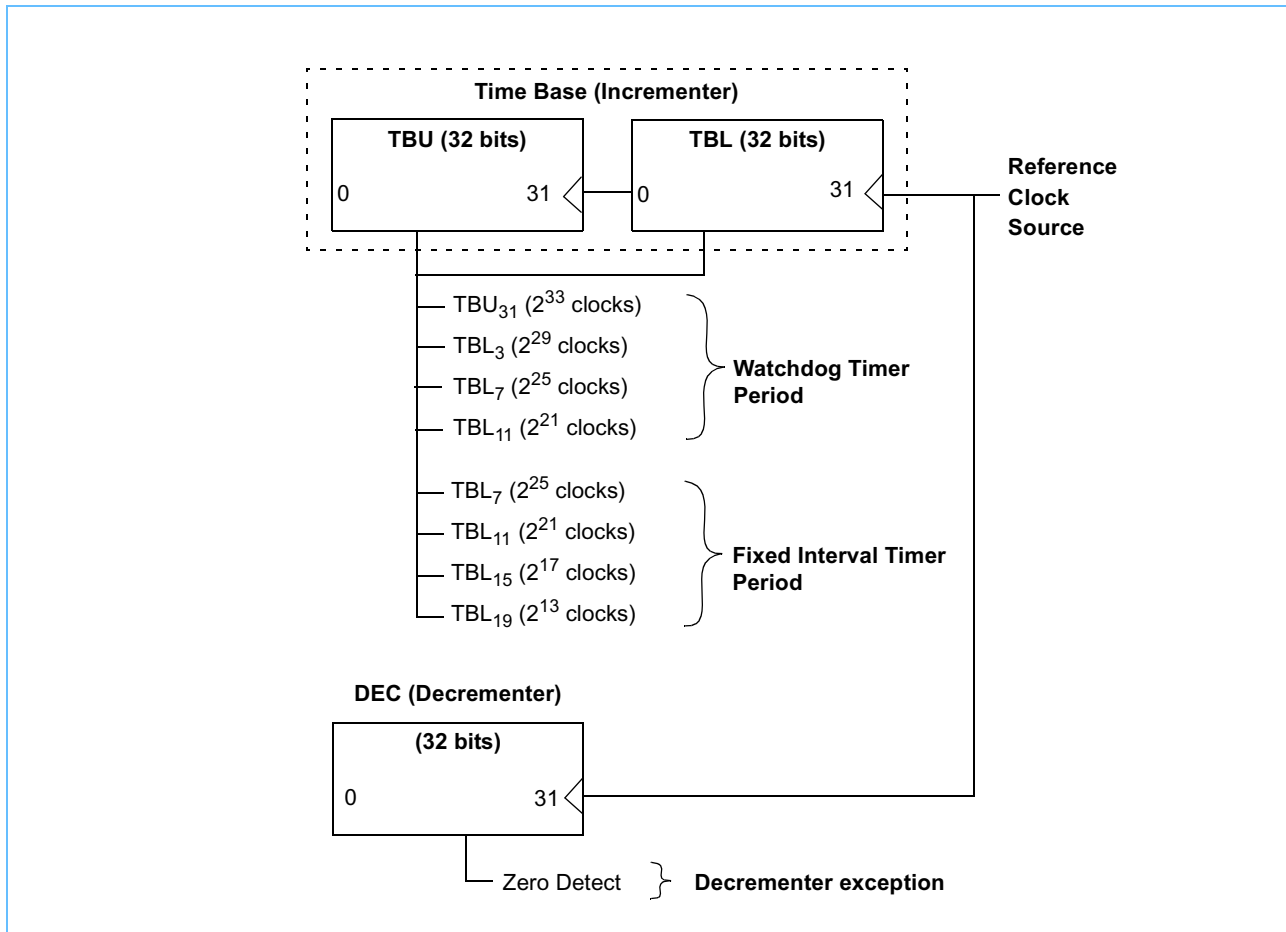
## 7. Timer Facilities

The PPC440x5 provides four timer facilities: a time base, a Decrementer (DEC), a Fixed Interval Timer (FIT), and a Watchdog Timer. These facilities, which share the same source clock frequency, can support:

- Time-of-day functions
- General software timing functions
- Peripherals requiring periodic service
- General system maintenance
- System error recoverability

Figure 7-1 shows the relationship between these facilities and the clock source.

Figure 7-1. Relationship of Timer Facilities to the Time Base



## 7.1 Time Base

The time base is a 64-bit register which increments once during each period of the source clock, and provides a time reference. Access to the time base is via two Special Purpose Registers (SPRs). The Time Base Upper (TBU) SPR contains the high-order 32 bits of the time base, while the Time Base Lower (TBL) SPR contains the low-order 32 bits.

Software access to TBU and TBL is non-privileged for read but privileged for write, and hence different SPR numbers are used for reading than for writing. TBU and TBL are written using **mtspr** and read using **mfspr**.

The period of the 64-bit time base is approximately 1462 years for a 400 MHz clock source. The time base value itself does not generate any exceptions, even when it wraps. For most applications, the time base is set once at system reset and only read thereafter. Note that Fixed Interval Timer and Watchdog Timer exceptions (discussed below) are caused by 0→1 transitions of selected bits from the time base. Transitions of these bits caused by software alteration of the time base have the same effect as transitions caused by normal incrementing of the time base.

Figure 7-2 illustrates the TBL.




Figure 7-2. Time Base Lower (TBL)



Figure 7-3 illustrates the TBU.




Figure 7-3. Time Base Upper (TBU)



### 7.1.1 Reading the Time Base

The following code provides an example of reading the time base.

```

loop:
    mfspr Rx,TBU          # read TBU into GPR Rx
    mfspr Ry,TBL          # read TBL into GPR Ry
    mfspr Rz,TBU          # read TBU again, this time into GPR Rz
    cmpw Rz, Rx           # see if old = new
    bne  loop             # loop/reread if rollover occurred

```

The comparison and loop ensure that a consistent pair of values is obtained.

### 7.1.2 Writing the Time Base

The following code provides an example of writing the time base.

```

lwz   Rx, upper          # load 64-bit time base value into GPRs Rx and Ry
lwz   Ry, lower
li    Rz, 0              # set GPR Rz to 0
mtspr TBL,Rz            # force TBL to 0 (thereby preventing wrap into TBU)
mtspr TBU,Rx            # set TBU to initial value
mtspr TBL,Ry            # set TBL to initial value

```

## 7.2 Decrementer (DEC)

The DEC is a 32-bit privileged SPR that decrements at the same rate that the time base increments. The DEC is read and written using **mfspr** and **mtspr**, respectively. When a non-zero value is written to the DEC, it begins to decrement with the next time base clock. A Decrementer exception is signalled when a decrement occurs on a DEC count of 1, and the Decrementer Interrupt Status field of the Timer Status Register (TSR[DIS]; see page 210) is set. A Decrementer interrupt will occur if it is enabled by both the Decrementer Interrupt Enable field of the Timer Control Register (TCR[DIE]; see page 209) and by the External Interrupt Enable field of the Machine State Register (MSR[EE]; see *Machine State Register (MSR)* on page 159). *Interrupts and Exceptions* on page 153 provides more information on the handling of Decrementer interrupts.

The Decrementer interrupt handler software should clear TSR[DIS] before re-enabling MSR[EE], in order to avoid another Decrementer interrupt due to the same exception (unless TCR[DIE] is cleared instead).

The behavior of the DEC itself upon a decrement from a DEC value of 1 depends on which of two modes it is operating in -- normal, or auto-reload. The mode is controlled by the Auto-Reload Enable (ARE) field of the TCR. When operating in normal mode (TCR[ARE]=0), the DEC simply decrements to the value 0 and then stops decrementing until it is re-initialized by software.

When operating in auto-reload mode (TCR[ARE]=1), however, instead of decrementing to the value 0, the DEC is reloaded with the value in the Decrementer Auto-Reload (DECAR) register (see *Figure 7-5*), and continues to decrement with the next time base clock (assuming the DECAR value was non-zero). The DECAR register is a 32-bit privileged, write-only SPR, and is written using **mtspr**.

The auto-reload feature of the DEC is disabled upon reset, and must be enabled by software.

*Figure 7-4* illustrates the DEC.




Figure 7-4. Decrementer (DEC)




Figure 7-5 illustrates the DECAR.





Figure 7-5. Decrementer Auto-Reload (DECAR)



Using **mtspr** to force the DEC to 0 does *not* cause a Decrementer exception and thus does not cause TSR[DIS] to be set. However, if a time base clock causes a decrement from a DEC value of 1 to occur simultaneously with the writing of the DEC by a **mtspr** instruction, then the Decrementer exception *will* occur, TSR[DIS] will be set, and the DEC will be written with the value from the **mtspr**.

In order for software to quiesce the activity of the DEC and eliminate all DEC exceptions, the following procedure should be followed:

1. Write 0 to TCR[DIE]. This prevents a Decrementer exception from causing a Decrementer interrupt.
2. Write 0 to TCR[ARE]. This disables the DEC auto-reload feature.
3. Write 0 to the DEC to halt decrementing. Although this action does not itself cause a Decrementer exception, it is possible that a decrement from a DEC value of 1 has occurred since the last time that TSR[DIS] was cleared.
4. Write 1 to TSR[DIS] (DEC Interrupt Status bit). This clears the Decrementer exception by setting TSR[DIS] to 0. Because the DEC is no longer decrementing (due to having been written with 0 in step 3), no further Decrementer exceptions are possible.

### 7.3 Fixed Interval Timer (FIT)

The FIT provides a mechanism for causing periodic exceptions with a regular period. The FIT would typically be used by system software to invoke a periodic system maintenance function, executed by the Fixed Interval Timer interrupt handler.

A Fixed Interval Timer exception occurs on a 0→1 transition of a selected bit from the time base. Note that a Fixed Interval Timer exception will also occur if the selected time base bit transitions from 0→1 due to a **mtspr** instruction that writes 1 to that time base bit when its previous value was 0.

The Fixed Interval Timer Period (FP) field of the TCR selects one of four bits from the time base, as shown in *Table 7-1*.

*Table 7-1. Fixed Interval Timer Period Selection*

TCR[FP]	Time Base Bit	Period (Time Base Clocks)	Period (400 Mhz Clock)
0b00	TBL <sub>19</sub>	2 <sup>13</sup> clocks	20.48 μs
0b01	TBL <sub>15</sub>	2 <sup>17</sup> clocks	327.68 μs
0b10	TBL <sub>11</sub>	2 <sup>21</sup> clocks	5.2 ms
0b11	TBL <sub>7</sub>	2 <sup>25</sup> clocks	83.9 ms

When a Fixed Interval Timer exception occurs, the exception status is recorded by setting the Fixed interval Timer Interrupt Status (FIS) field of the TSR to 1. A Fixed Interval Timer interrupt will occur if it is enabled by both the Fixed Interval Timer Interrupt Enable (FIE) field of the TCR and by MSR[EE]. *Fixed-Interval Timer Interrupt* on page 185 provides more information on the handling of Fixed Interval Timer interrupts.

The Fixed Interval Timer interrupt handler software should clear TSR[FIS] before re-enabling MSR[EE], in order to avoid another Fixed Interval Timer interrupt due to the same exception (unless TCR[FIE] is cleared instead).

## 7.4 Watchdog Timer

The Watchdog Timer provides a mechanism for system error recovery in the event that the program running on the PPC440x5 core has stalled and cannot be interrupted by the normal interrupt mechanism. The Watchdog Timer can be configured to cause a critical-class Watchdog Timer interrupt upon the expiration of a single period of the Watchdog Timer. It can also be configured to invoke a processor-initiated reset upon the expiration of a second period of the Watchdog Timer.

A Watchdog Timer exception occurs on a 0→1 transition of a selected bit from the time base. Note that a Watchdog Timer exception will also occur if the selected time base bit transitions from 0→1 due to a **mtspr** instruction that writes 1 to that time base bit when its previous value was 0.

The Watchdog Timer Period (WP) field of the TCR selects one of four bits from the time base, as shown in *Table 7-2*.

*Table 7-2. Watchdog Timer Period Selection*

TCR[WP]	Time Base Bit	Period (Time Base Clocks)	Period (400 MHz Clock)
0b00	TBL <sub>11</sub>	2 <sup>21</sup> clocks	5.2 ms
0b01	TBL <sub>7</sub>	2 <sup>25</sup> clocks	83.9 ms
0b10	TBL <sub>3</sub>	2 <sup>29</sup> clocks	1.34 s
0b11	TBU <sub>31</sub>	2 <sup>33</sup> clocks	21.47 s

The action taken upon a Watchdog Timer exception depends upon the status of the Enable Next Watchdog (ENW) and Watchdog Timer Interrupt Status (WIS) fields of the TSR at the time of the exception. When  $TSR[ENW] = 0$ , the next Watchdog Timer exception is “disabled”, and the only action to be taken upon the exception is to set  $TSR[ENW]$  to 1. By clearing  $TSR[ENW]$ , software can guarantee that the time until the next *enabled* Watchdog Timer exception will be *at least* one full Watchdog Timer period (and a maximum of *two* full Watchdog Timer periods).

When  $TSR[ENW] = 1$ , the next Watchdog Timer exception is enabled, and the action to be taken upon the exception depends on the value of  $TSR[WIS]$  at the time of the exception. If  $TSR[WIS] = 0$ , then the action is to set  $TSR[WIS]$  to 1, at which time a Watchdog Timer interrupt will occur if enabled by both the Watchdog Timer Interrupt Enable (WIE) field of the TCR and by the Critical Interrupt Enable (CE) field of the MSR. The Watchdog Timer interrupt handler software should clear  $TSR[WIS]$  before re-enabling  $MSR[CE]$ , in order to avoid another Watchdog Timer interrupt due to the same exception (unless  $TCR[WIE]$  is cleared instead). *Watchdog Timer Interrupt* on page 186 provides more information on the handling of Watchdog Timer interrupts.

If  $TSR[WIS]$  is already 1 at the time of the next Watchdog Timer exception, then the action to take depends on the value of the Watchdog Reset Control (TRC) field of the TCR. If  $TCR[WRC]$  is non-zero, then the value of the  $TCR[WRC]$  field will be copied into  $TSR[WRS]$ ,  $TCR[WRC]$  will be cleared, and a core reset will occur (see *PPC440x5 Core State After Reset* on page 83 for more information on core behavior when reset).

Note that once software has set  $TCR[WRC]$  to a non-zero value, it cannot be reset by software; this feature prevents errant software from disabling the Watchdog Timer reset capability.

Table 7-3 summarizes the action to be taken upon a Watchdog Timer exception according to the values of  $TSR[ENW]$  and  $TSR[WIS]$ .

Table 7-3. Watchdog Timer Exception Behavior

TSR[ENW]	TSR[WIS]	Action upon Watchdog Timer exception
0	0	Set $TSR[ENW]$ to 1
0	1	Set $TSR[ENW]$ to 1
1	0	Set $TSR[WIS]$ to 1. If Watchdog Timer interrupts are enabled ( $TCR[WIE]=1$ and $MSR[CE]=1$ ), then interrupt.
1	1	Cause Watchdog Timer reset action specified by $TCR[WRC]$ . Reset will copy pre-reset $TCR[WRC]$ into $TSR[WRS]$ , then clear $TCR[WRC]$ .

A typical system usage of the Watchdog Timer function is to enable the Watchdog Timer interrupt and the Watchdog Timer reset function in the TCR (and MSR), and to start out with both  $TSR[ENW]$  and  $TSR[WIS]$  clear. Then, a recurring software loop of reliable duration (or perhaps the interrupt handler for a periodic interrupt such as the Fixed Interval Timer interrupt) performs a periodic check of system integrity. Upon successful completion of the system check, software clears  $TSR[ENW]$ , thereby ensuring that a minimum of one full Watchdog Timer period and a maximum of two full Watchdog Timer periods must expire before an enabled Watchdog Timer exception will occur.

If for some reason the recurring software loop is not successfully completed (and  $TSR[ENW]$  thus not cleared) during this period of time, then an enabled Watchdog Timer exception will occur. This will set  $TSR[WIS]$  and a Watchdog Timer interrupt will occur (if enabled by both  $TCR[WIE]$  and  $MSR[CE]$ ). The occurrence of a Watchdog Timer interrupt in this kind of system is interpreted as a “system error”, insofar as the system was for some reason unable to complete the periodic system integrity check in time to avoid the enabled Watchdog Timer exception. The action taken by the Watchdog Timer interrupt handler is of course system-dependent, but typically the software will attempt to determine the nature of the problem and correct it

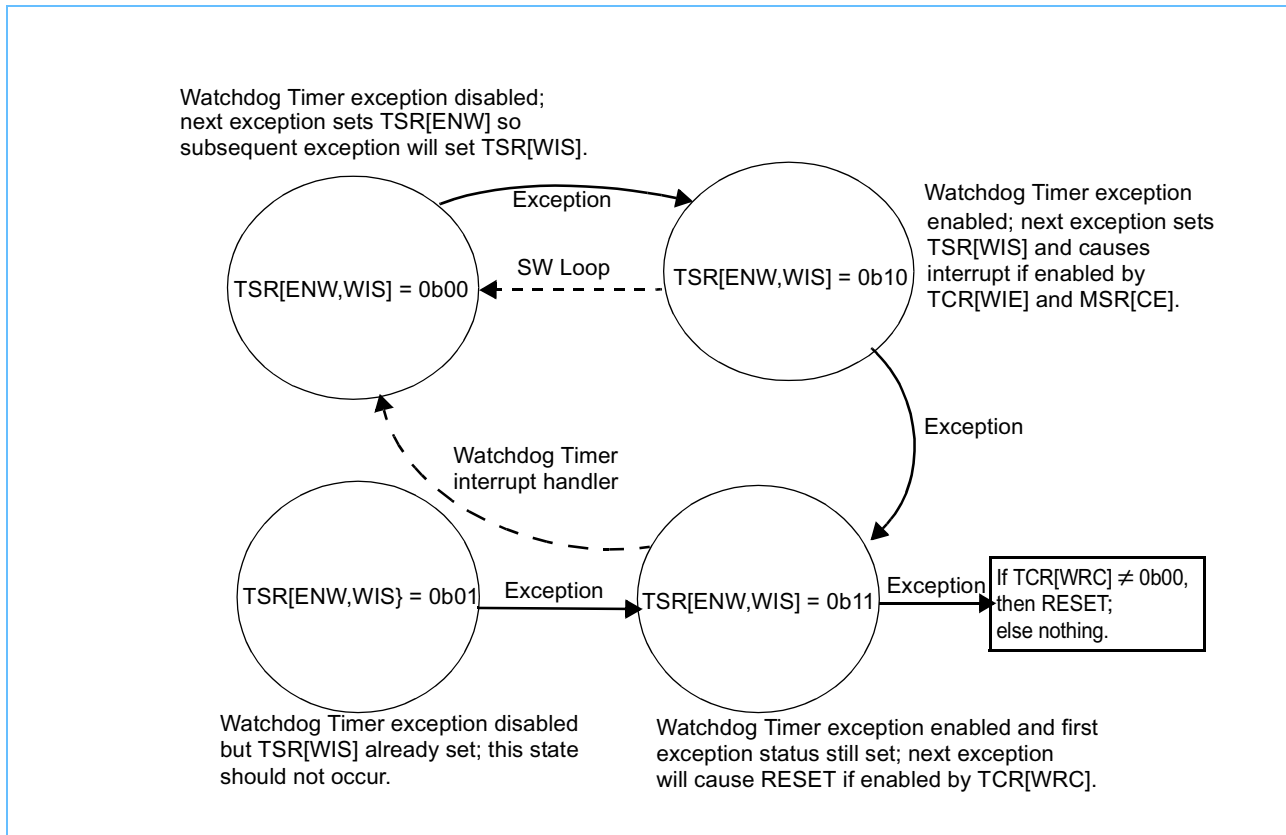


if possible. If and when the system attempts to resume operation, the software would typically clear both TSR[WIS] and TSR[ENW], thus providing a minimum of another full Watchdog Timer period for a new system integrity check to occur.

Finally, if for some reason the Watchdog Timer interrupt is disabled, and/or the Watchdog Timer interrupt handler is unsuccessful in clearing TSR[WIS] and TSR[ENW] prior to another Watchdog Timer exception, then the next exception will cause a processor reset operation to occur, according to the value of TCR[WRC].

Figure 7-6 illustrates the sequence of Watchdog Timer events which occurs according to this typical system usage.

Figure 7-6. Watchdog State Machine



## 7.5 Timer Control Register (TCR)

The TCR is a privileged SPR that controls DEC, FIT, and Watchdog Timer operation. The TCR is read into a GPR using `mfspr`, and is written from a GPR using `mtspr`.

The Watchdog Timer Reset Control (WRC) field of the TCR is cleared to 0 by processor reset (see *Initialization* on page 83). Each bit of this 2-bit field is set only by software and is cleared only by hardware. For each bit of the field, once software has written it to 1, that bit remains 1 until processor reset occurs. This is to prevent errant code from disabling the Watchdog Timer reset function.

The Auto-Reload Enable (ARE) field of the TCR is also cleared to zero by processor reset. This disables the auto-reload feature of the DEC.




Figure 7-7. Timer Control Register (TCR)

0:1	WP	Watchdog Timer Period 00 $2^{21}$ time base clocks 01 $2^{25}$ time base clocks 10 $2^{29}$ time base clocks 11 $2^{33}$ time base clocks	
2:3	WRC	Watchdog Timer Reset Control 00 No Watchdog Timer reset will occur. 01 Core reset 10 Chip reset 11 System reset	TCR[WRC] resets to 0b00. Type of reset to cause upon Watchdog Timer exception with TSR[ENW,WIS]=0b11. This field can be set by software, but cannot be cleared by software, except by a software-induced reset.
4	WIE	Watchdog Timer Interrupt Enable 0 Disable Watchdog Timer interrupt. 1 Enable Watchdog Timer interrupt.	
5	DIE	Decrementer Interrupt Enable 0 Disable Decrementer interrupt. 1 Enable Decrementer interrupt.	
6:7	FP	Fixed Interval Timer (FIT) Period 00 $2^{13}$ time base clocks 01 $2^{17}$ time base clocks 10 $2^{21}$ time base clocks 11 $2^{25}$ time base clocks	
8	FIE	FIT Interrupt Enable 0 Disable Fixed Interval Timer interrupt. 1 Enable Fixed Interval Timer interrupt.	
9	ARE	Auto-Reload Enable 0 Disable auto reload. 1 Enable auto reload.	TCR[ARE] resets to 0b0.
10:31		Reserved	

## 7.6 Timer Status Register (TSR)

The TSR is a privileged SPR that records the status of DEC, FIT, and Watchdog Timer events. The fields of the TSR are generally set to 1 only by hardware and cleared to 0 only by software. Hardware cannot clear any fields in the TSR, nor can software set any fields. Software can read the TSR into a GPR using **mfspr**. Clearing the TSR is performed using **mtspr** by placing a 1 in the GPR source register in all bit positions which are to be cleared in the TSR, and a 0 in all other bit positions. The data written from the GPR to the TSR is not direct data, but a mask. A 1 clears the bit and a 0 leaves the corresponding TSR bit unchanged.




Figure 7-8. Timer Status Register (TSR)

0	ENW	Enable Next Watchdog Timer Exception 0 Action on next Watchdog Timer exception is to set TSR[ENW] = 1. 1 Action on next Watchdog Timer exception is governed by TSR[WIS].
1	WIS	Watchdog Timer Interrupt Status 0 Watchdog Timer exception has not occurred. 1 Watchdog Timer exception has occurred.
2:3	WRS	Watchdog Timer Reset Status 00 No Watchdog Timer reset has occurred. 01 Core reset was forced by Watchdog Timer. 10 Chip reset was forced by Watchdog Timer. 11 System reset was forced by Watchdog Timer.
4	DIS	Decrementer Interrupt Status 0 Decrementer exception has not occurred. 1 Decrementer exception has occurred.
5	FIS	Fixed Interval Timer (FIT) Interrupt Status 0 Fixed Interval Timer exception has not occurred. 1 Fixed Interval Timer exception has occurred.
6:31		Reserved

## 7.7 Freezing the Timer Facilities

The debug mechanism provides a means for temporarily “freezing” the timers upon a debug exception. Specifically, the time base and Decrementer can be prevented from incrementing and decrementing, respectively, whenever a debug exception is recorded in the Debug Status Register (DBSR). This allows a debugger to simulate the appearance of “real time”, even though the application has been temporarily “halted” to service the debug event.

See *Debug Facilities* on page 213 for more information on freezing the timers.

## 7.8 Selection of the Timer Clock Source

The source clock of the timers is selected by the Timer Clock Select (TCS) field of the Core Configuration Register 1 (CCR1). When set to zero, CCR1[TCS] selects the CPU clock. This is the highest frequency timer clock source.

When set to one, CCR1[TCS] selects an input to the CPU core as the timer clock. The input is sampled by a latch clocked by the CPU clock, and so cannot cycle any faster than 1/2 the frequency of the CPU clock.



## 8. Debug Facilities

The debug facilities of the PPC440x5 include support for several debug modes for debugging during hardware and software development, as well as debug events that allow developers to control the debug process. Debug registers control these debug modes and debug events. The debug registers may be accessed either through software running on the processor or through the JTAG debug port of the PPC440x5 core. Access to the debug facilities through the JTAG debug port is typically provided by a debug tool such as the RISC-Watch™ development tool from IBM. A trace port, which enables the tracing of code running in real time, is also provided.

### 8.1 Support for Development Tools

The RISCWatch product from IBM is an example of a development tool that uses external debug mode, debug events, and the JTAG debug port to implement a hardware and software development tool. The RISC-Trace™ feature of RISCWatch is an example of a development tool that uses the real-time instruction trace capability of the PPC440x5 core.

### 8.2 Debug Modes

The following sections describe the various debug modes supported by the PPC440x5. Each of these debug modes supports a particular type of debug tool or debug task commonly used in embedded systems development. For all debug modes, the various debug event types are enabled by the setting of corresponding fields in Debug Control Register 0 (DBCR0), and upon their occurrence are recorded in the Debug Status Register (DBSR).

There are four debug modes:

- Internal debug mode
- External debug mode
- Debug wait mode
- Trace debug mode

The PowerPC Book-E architecture specification deals only with internal debug mode, and the relationship of Debug interrupts to the rest of the interrupt architecture. Internal debug mode is the mode which involves debug software running on the processor itself, typically in the form of the Debug interrupt handler. The other debug modes, on the other hand, are outside the scope of the architecture, and involve special-purpose debug hardware external to the PPC440x5 core, connected either to the JTAG interface (for external debug mode and debug wait mode) or the trace interface (for trace debug mode). Details of these interfaces and their operation are beyond the scope of this manual.

#### 8.2.1 Internal Debug Mode

Internal debug mode provides access to architected processor resources and supports setting hardware and software breakpoints and monitoring processor status. In this mode, debug events are considered exceptions, which, in addition to recording their status in the DBSR, generate Debug interrupts if and when such interrupts are enabled (Machine State Register (MSR) DE field is 1; see *Interrupts and Exceptions* on

page 153 for a description of the MSR and Debug interrupts). When a Debug interrupt occurs, special debugger software at the interrupt handler can check processor status and other conditions related to the debug event, as well as alter processor resources using all of the instructions defined for the PPC440x5.

Internal debug mode relies on this interrupt handling software at the Debug interrupt vector to debug software problems. This mode, used while the processor executes instructions, enables debugging of both application programs and operating system software, including all of the non-critical class interrupt handlers.

In this mode, the debugger software can communicate with the outside world through a communications port, such as a serial port, external to the processor core.

To enable internal debug mode, the IDM field of DBCR0 must be set to 1 (DBCR0[IDM] = 1). This mode can be enabled in combination with external debug mode (see *External Debug Mode* below) and/or debug wait mode (see *Debug Wait Mode* on page 214).

### 8.2.2 External Debug Mode

External debug mode provides access to architected processor resources and supports stopping, starting, and stepping the processor, setting hardware and software breakpoints, and monitoring processor status. In this mode, debug events record their status in the DBSR and then cause the processor to enter the *stop state*, in which normal instruction execution stops and architected processor resources and memory can be accessed and altered via the JTAG interface. While in the stop state, interrupts are temporarily disabled.

Storage access control by a memory management unit (MMU) remains in effect while in external debug mode; the debugger may need to modify MSR or TLB values to access protected memory.

External debug mode relies only on internal processor resources, and no Debug interrupt handling software, so it can be used to debug both system hardware and software problems. This mode can also be used for software development on systems without a control program, or to debug control program problems, including problems within the Debug interrupt handler itself, or within any other critical class interrupt handlers.

External debug mode is enabled by setting DBCR0[EDM] to 1. This mode can be enabled in combination with internal debug mode (see *Internal Debug Mode* on page 213) and/or debug wait mode (see *Debug Wait Mode* below). External debug mode takes precedence over internal debug mode however, in that debug events will first cause the processor to enter stop state rather than generating a Debug interrupt, although a Debug interrupt may be pending while the processor is in the stop state.

### 8.2.3 Debug Wait Mode

Debug wait mode is similar to external debug mode in that debug events cause the processor to enter the stop state. However, interrupts are still enabled while in debug wait mode, such that if and when an exception occurs for which the associated interrupt type is enabled, the processor will leave the stop state and generate the interrupt. This mode is useful for real-time hardware environments which cannot tolerate interrupts being disabled for an extended period of time. In such environments, if external debug mode were to be used, various I/O devices could operate incorrectly due to not being serviced in a timely fashion when they assert an interrupt request to the processor, if the processor happened to be in stop state at the time of the interrupt request.

When in debug wait mode, as with external debug mode, access to the architected processor resources and memory is via the JTAG interface.

Debug wait mode is enabled by setting both MSR[DWE] and the debug wait mode enable within the JTAG controller to 1. Since MSR[DWE] is automatically cleared upon any interrupt, debug wait mode is temporarily disabled upon an interrupt, and then can be automatically re-enabled when returning from the interrupt due to the restoration of the MSR value upon the execution of an **rfi**, **rftci**, or **rfmci** instruction.

While debug wait mode can be enabled in combination with external debug mode, external debug mode takes precedence and interrupts are temporarily disabled, thereby effectively nullifying the effect of debug wait mode. Similarly, debug wait mode can be enabled in combination with internal debug mode. However, if Debug interrupts are enabled (MSR[DE] is 1), then any debug event will lead to an exception and a corresponding Debug interrupt, which takes precedence over the stop state associated with debug wait mode. On the other hand, if Debug interrupts are disabled (MSR[DE] is 0), then debug wait mode will take effect and a debug event will cause the processor to enter stop state.

### 8.2.4 Trace Debug Mode

Trace debug mode is simply the *absence* of each of the other modes. That is, if internal debug mode, external debug mode, and debug wait mode are all disabled, then the processor is in trace debug mode. While in trace debug mode, all debug events are simply recorded in the DBSR, and are indicated over the trace interface from the PPC440x5 core. The processor does not enter the stop state, nor does a Debug interrupt occur.

## 8.3 Debug Events

There are several different kinds of debug events, each of which is enabled by a field in DBCR0 (except for the Unconditional debug event) and recorded in the DBSR. *Debug Modes* on page 213 describes the operation that results when a debug event occurs while operating in any of the debug modes.

Table 8-1 lists the various debug events recognized by the PPC440x5. Detailed explanations of each debug event type follow the table.

Table 8-1. Debug Events

Event	Description
Instruction Address Compare (IAC)	Caused by the attempted execution of an instruction for which the address matches the conditions specified by DBCR0, DBCR1, and the IAC1–IAC4 registers.
Data Address Compare (DAC)	Caused by the attempted execution of a load, store, or cache management instruction for which the data storage address matches the conditions specified by DBCR0, DBCR2, and the DAC1–DAC2 registers.
Data Value Compare (DVC)	Caused by the attempted execution of a load, store, or cache management instruction for which the data storage address matches the conditions specified by DBCR0, DBCR2, and the DAC1–DAC2 registers, and for which the referenced data matches the value specified by the DVC1–DVC2 registers.
Branch Taken (BRT)	Caused by the attempted execution of a branch instruction for which the branch conditions are met (that is, for a branch instruction that results in the re-direction of the instruction stream).
Trap (TRAP)	Caused by the attempted execution of a <b>tw</b> or <b>twi</b> instruction for which the trap conditions are met.
Return (RET)	Caused by the attempted execution of an <b>rftci</b> , <b>rftci</b> , or <b>rfmci</b> instruction.

Table 8-1. Debug Events

Event	Description
Instruction Complete (ICMP)	Caused by the successful completion of the execution of any instruction.
Interrupt (IRPT)	Caused by the generation of an interrupt.
Unconditional (UDE)	Caused by the assertion of an unconditional debug event request from the JTAG interface to the PPC440x5 core.

### 8.3.1 Instruction Address Compare (IAC) Debug Event

IAC debug events occur when execution is attempted of an instruction for which the instruction address and other parameters match the IAC conditions specified by DBCR0, DBCR1, and the IAC registers. There are four IAC registers on the PPC440x5, IAC1–IAC4. Depending on the IAC mode specified by DBCR1, these IAC registers can be used to specify four independent, exact IAC addresses, or they can be configured in pairs (IAC1/IAC2 and IAC3/IAC4) in order to specify *ranges* of instruction addresses for which IAC debug events should occur.

#### 8.3.1.1 IAC Debug Event Fields

There are several fields in DBCR0 and DBCR1 which are used to specify the IAC conditions, as follows:

##### IAC Event Enable Field

DBCR0[IAC1, IAC2, IAC3, IAC4] are the individual IAC event enables for each of the four IAC events: IAC1, IAC2, IAC3, and IAC4. For a given IAC event to occur, the corresponding IAC event enable bit in DBCR0 must be set. When a given IAC event occurs, the corresponding DBSR[IAC1, IAC2, IAC3, IAC4] bit is set.

##### IAC Mode Field

DBCR1[IAC12M, IAC34M] control the comparison mode for the IAC1/IAC2 and IAC3/IAC4 events, respectively. There are three comparison modes supported by the PPC440x5:

- Exact comparison mode (DBCR1[IAC12M/IAC34M] = 0b00)

In this mode, the instruction address is compared to the value in the corresponding IAC register, and the IAC event occurs only if the comparison is an exact match.

- Range inclusive comparison mode (DBCR1[IAC12M/IAC34M] = 0b10)

In this mode, the IAC1 or IAC2 event occurs only if the instruction address is *within* the range defined by the IAC1 and IAC2 register values, as follows:  $IAC1 \leq \text{address} < IAC2$ . Similarly, the IAC3 or IAC4 event occurs only if the instruction address is *within* the range defined by the IAC3 and IAC4 register values, as follows:  $IAC3 \leq \text{address} < IAC4$ .

For a given IAC1/IAC2 or IAC3/IAC4 pair, when the instruction address falls within the specified range, either one or both of the corresponding IAC debug event bits will be set in the DBSR, as determined by which of the two corresponding IAC event enable bits are set in DBCR0. For example, when the IAC1/IAC2 pair are set to range inclusive comparison mode, and the instruction address falls within the defined range, then DBCR1[IAC1, IAC2] will determine whether one or the other or both of DBSR[IAC1, IAC2] are set. It is a programming error to set either of the IAC pairs to a range



comparison mode (either inclusive or exclusive) without also enabling at least one of the corresponding IAC event enable bits in DBCR0.

Note that the IAC range auto-toggle mechanism can “switch” the IAC range mode from inclusive to exclusive, and vice-versa. See *IAC Range Mode Auto-Toggle Field* on page 218.

- Range exclusive comparison mode (DBCR1[IAC12M/IAC34M] = 0b11)

In this mode, the IAC1 or IAC2 event occurs only if the instruction address is *outside* the range defined by the IAC1 and IAC2 register values, as follows: address < IAC1 or address ≥ IAC2. Similarly, the IAC3 or IAC4 event occurs only if the instruction address is *outside* the range defined by the IAC3 and IAC4 register values, as follows: address < IAC3 or address ≥ IAC4.

For a given IAC1/IAC2 or IAC3/IAC4 pair, when the instruction address falls outside the specified range, either one or both of the corresponding IAC debug event bits will be set in the DBSR, as determined by which of the two corresponding IAC event enable bits are set in DBCR0. For example, when the IAC1/IAC2 pair are set to range exclusive comparison mode, and the instruction address falls outside the defined range, then DBCR1[IAC1, IAC2] will determine whether one or the other or both of DBSR[IAC1, IAC2] are set. It is a programming error to set either of the IAC pairs to a range comparison mode (either inclusive or exclusive) without also enabling at least one of the corresponding IAC event enable bits in DBCR0.

Note that the IAC range auto-toggle mechanism can “switch” the IAC range mode from inclusive to exclusive, and vice-versa. See *IAC Range Mode Auto-Toggle Field* on page 218.

The PowerPC Book-E architecture defines DBCR1[IAC12M/IAC34M] = 0b01 as IAC address bit mask mode, but that mode is not supported by the PPC440x5, and that value of the IAC12M/IAC34M fields is reserved.

### IAC User/Supervisor Field

DBCR1[IAC1US, IAC2US, IAC3US, IAC4US] are the individual IAC user/supervisor fields for each of the four IAC events. The IAC user/supervisor fields specify what operating mode the processor must be in order for the corresponding IAC event to occur. The operating mode is determined by the Problem State field of the Machine State Register (MSR[PR]; see *User and Supervisor Modes* on page 78). When the IAC user/supervisor field is 0b00, the operating mode does not matter; the IAC debug event may occur independent of the state of MSR[PR]. When this field is 0b10, the processor must be operating in supervisor mode (MSR[PR] = 0). When this field is 0b11, the processor must be operating in user mode (MSR[PR] = 1). The IAC user/supervisor field value of 0b01 is reserved.

If a pair of IAC events (IAC1/IAC2 or IAC3/IAC4) are operating in range inclusive or range exclusive mode, it is a programming error (and the results of any instruction address comparison are undefined) if the corresponding pair of IAC user/supervisor fields are not set to the same value. For example, if IAC1/IAC2 are operating in one of the range modes, then both DBCR1[IAC1US] and DBCR1[IAC2US] must be set to the same value.

### IAC Effective/Real Address Field

DBCR1[IAC1ER, IAC2ER, IAC3ER, IAC4ER] are the individual IAC effective/real address fields for each of the four IAC events. The IAC effective/real address fields specify whether the instruction address comparison should be performed using the effective, virtual, or real address (see *Memory Management* on page 129) for an explanation of these different types of addresses). When the IAC effective/real address field is 0b00, the comparison is performed using the effective address only—the IAC debug event may

occur independent of the instruction address space (MSR[IS]). When this field is 0b10, the IAC debug event occurs only if the effective address matches the IAC conditions and is in virtual address space 0 (MSR[IS] = 0). Similarly, when this field is 0b11, the IAC debug event occurs only if the effective address matches the IAC conditions and is in virtual address space 1 (MSR[IS] = 1). Note that in these latter two modes, in which the virtual address space of the instruction is considered, it is not the entire virtual address which is considered. The Process ID, which forms the final part of the virtual address, is not considered. Finally, the IAC effective/real address field value of 0b01 is reserved, and corresponds to the PowerPC Book-E architected real address comparison mode, which is not supported by the PPC440x5.

If a pair of IAC events (IAC1/IAC2 or IAC3/IAC4) are operating in range inclusive or range exclusive mode, it is a programming error (and the results of any instruction address comparison are undefined) if the corresponding pair of IAC effective/real address fields are not set to the same value. For example, if IAC1/IAC2 are operating in one of the range modes, then both DBCR1[IAC1ER] and DBCR1[IAC2ER] must be set to the same value.

### IAC Range Mode Auto-Toggle Field

DBCR1[IAC12AT, IAC34AT] control the auto-toggle mechanism for the IAC1/IAC2 and IAC3/IAC4 events, respectively. When the IAC mode for one of the pairs of IAC debug events is set to one of the range modes (either range inclusive or range exclusive), then the IAC range mode auto-toggle field corresponding to that pair of IAC debug events controls whether or not the range mode will automatically “toggle” from inclusive to exclusive, and vice-versa. When the IAC range mode auto-toggle field is set to 1, this automatic toggling is enabled; otherwise it is disabled. It is a programming error (and the results of any instruction address comparison are undefined) if an IAC range mode auto-toggle field is set to 1 without the corresponding IAC mode field being set to one of the range modes.

When auto-toggle is enabled for a pair of IAC debug events, then upon each occurrence of an IAC debug event within that pair the value of the corresponding auto-toggle status field in the DBSR (DBSR[IAC12ATS, IAC34ATS]) is reversed. That is, if the auto-toggle status field is 0 before the occurrence of the IAC debug event, then it will be changed to 1 at the same time that the IAC debug event is recorded in the DBSR. Conversely, if the auto-toggle status field is 1 before the occurrence of the IAC debug event, then it will be changed to 0 at the same time that the IAC debug event is recorded in the DBSR.

Furthermore, when auto-toggle is enabled, the auto-toggle status field of the DBSR affects the interpretation of the IAC mode field of DBCR1. If the auto-toggle status field is 0, then the IAC mode field is interpreted in the normal fashion, as defined in *IAC Mode Field* on page 216. That is, the IAC mode field value of 0b10 selects range inclusive mode, whereas the value of 0b11 selects range exclusive mode. On the other hand, when the auto-toggle status field is 1, then the interpretation of the IAC mode field is “reversed”. That is, the IAC mode field value of 0b10 selects range exclusive mode, whereas the value of 0b11 selects range inclusive mode.

The relationship of the IAC mode, IAC range mode auto-toggle, and IAC range mode auto-toggle status fields is summarized in *Table 8-2*.

Table 8-2. IAC Range Mode Auto-Toggle Summary

DBCR1 IAC12M/IAC34M	DBCR1 IAC12AT/IAC34AT	DBSR IAC12ATS/IAC34ATS	IAC Mode
0b10	0	—	Range Inclusive
0b10	1	0	Range Inclusive
0b10	1	1	Range Exclusive
0b11	0	—	Range Exclusive
0b11	1	0	Range Exclusive
0b11	1	1	Range Inclusive

The affect of the auto-toggle mechanism is to cause the IAC mode to switch back and forth between range inclusive mode and range exclusive mode, as each IAC range mode debug event occurs. For example, if the IAC mode is set to range inclusive, and auto-toggle is enabled, and the auto-toggle status field is 0, then the first IAC debug event will be a range inclusive event. Upon that event, the DBSR auto-toggle status field will be set to 1, and the next IAC debug event will then be a range exclusive event. Upon this next event, the DBSR auto-toggle status field will be set back to 0, such that the next IAC debug event will again be a range inclusive event.

This auto-toggling between range inclusive and range exclusive IAC modes is particularly helpful when enabling IAC range mode debug events in trace debug mode. A common debug operation is to detect when the instruction stream enters a particular region of the instruction address space (range inclusive mode). Once having entered the region of interest (a range inclusive event), it is common for the debugger to then want to be informed when that region is exited (a range exclusive event). By automatically toggling to range exclusive mode upon the occurrence of the range inclusive IAC debug event, this particular debug operation is facilitated. Furthermore, by not remaining in range inclusive mode upon entry to the region of interest, the debugger avoids a continuous stream of range inclusive IAC debug events while the processor continues to execute instructions within that region, which can often be for a very long series of instructions.

### 8.3.1.2 IAC Debug Event Processing

When operating in external debug mode or debug wait mode, the occurrence of an IAC debug event is recorded in the corresponding bit of the DBSR and causes the instruction execution to be suppressed. The processor then enters the stop state and ceases the processing of instructions. The program counter will contain the address of the instruction which caused the IAC debug event. Similarly, when operating in internal debug mode with Debug interrupts enabled (MSR[DE] = 1), the occurrence of an IAC debug event is recorded in the DBSR and causes the instruction execution to be suppressed. A Debug interrupt then occurs with Critical Save/Restore Register 0 (CSRR0) set to the address of the instruction which caused the IAC debug event.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled (MSR[DE] = 0), the behavior of IAC debug events depends on the IAC mode. If the IAC mode is set to exact comparison, then an IAC debug event can occur and will set the corresponding IAC field of the DBSR, along with the Imprecise Debug Event (IDE) field of the DBSR. The instruction execution is not suppressed, as no Debug interrupt will occur immediately. Instead, instruction execution continues, and a Debug interrupt will occur if and when MSR[DE] is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the IAC debug event status from the DBSR in the meantime. Upon such a “delayed”

interrupt, the Debug interrupt handler software may query the DBSR[IDE] field to determine that the Debug interrupt has occurred imprecisely. On the other hand, if the IAC mode is set to either range inclusive or range exclusive mode, then IAC debug events cannot occur when operating in internal debug mode with MSR[DE] = 0, unless external debug mode and/or debug wait mode is also enabled.

When operating in trace mode, the occurrence of an IAC debug event simply sets the corresponding IAC field of the DBSR and is indicated over the trace interface, and instruction execution continues.

### 8.3.2 Data Address Compare (DAC) Debug Event

DAC debug events occur when execution is attempted of a load, store, or cache management instruction for which the data storage address and other parameters match the DAC conditions specified by DBCR0, DBCR2, and the DAC registers. There are two DAC registers on the PPC440x5, DAC1 and DAC2. Depending on the DAC mode specified by DBCR2, these DAC registers can be used to specify two independent, exact DAC addresses, or they can be configured to operate as a pair. When operating as a pair, then can specify either a *range* of data storage addresses for which DAC debug events should occur, or a combination of an address and an *address bit mask* for selective comparison with the data storage address.

Note that for integer load and store instructions, and for cache management instructions, the address that is used in the DAC comparison is the starting data address calculated as part of the instruction execution. As explained in the instruction definitions for the cache management instructions, the target operand of these instructions is an aligned cache block, which on the PPC440x5 is 32 bytes. Therefore, the storage reference for these instructions effectively ignores the low-order five bits of the calculated data address, and the entire aligned 32-byte cache block—which starts at the calculated data address as modified with the low-order five bits set to 0b00000—is accessed. However, the DAC comparison does not take into account this implicit 32-byte alignment of the storage reference of a cache management instruction, and instead the DAC comparison considers the entire data address, as calculated according to the instruction definition.

On the other hand, for auxiliary processor load and store instructions, the AP interface can specify that the PPC440x5 should *force* the storage access to be aligned on an operand-size boundary, by zeroing the appropriate number of low-order address bits. In such a case, the DAC comparison is performed against this modified, alignment-forced address, rather than the original address as calculated according to the instruction definition.

#### 8.3.2.1 DAC Debug Event Fields

There are several fields in DBCR0 and DBCR2 which are used to specify the DAC conditions, as follows:

##### DAC Event Enable Field

DBCR0[*DAC1R*, *DAC1W*, *DAC2R*, *DAC2W*] are the individual DAC event enables for the two DAC events DAC1 and DAC2. For each of the two DAC events, there is one enable for DAC read events, and another for DAC write events. Load, **dcbt**, **dcbtst**, **icbi**, and **icbt** instructions may cause DAC read events, while store, **dcbst**, **dcbf**, **dcbi**, and **dcbz** instructions may cause DAC write events (see *DAC Debug Events Applied to Various Instruction Types* on page 224 for more information on these instructions and the types of DAC debug events they may cause). For a given DAC event to occur, the corresponding DAC event enable bit in DBCR0 for the particular operation type must be set. When a DAC event occurs, the corresponding DBSR[*DAC1R*, *DAC1W*, *DAC2R*, *DAC2W*] bit is set. These same DBSR bits are shared by DVC debug events (see *Data Value Compare (DVC) Debug Event* on page 225).

## DAC Mode Field

DBCR2[**DAC12M**] controls the comparison mode for the DAC1 and DAC2 events. There are four comparison modes supported by the PPC440x5:

- Exact comparison mode (DBCR2[**DAC12M**] = 0b00)

In this mode, the data address is compared to the value in the corresponding DAC register, and the DAC event occurs only if the comparison is an exact match.

- Address bit mask mode (DBCR2[**DAC12M**] = 0b01)

In this mode, the DAC1 or DAC2 event occurs only if the data address matches the value in the DAC1 register, as masked by the value in the DAC2 register. That is, the DAC1 register specifies an address value, and the DAC2 register specifies an *address bit mask* which determines which bit of the data address should participate in the comparison to the DAC1 value. For every bit set to 1 in the DAC2 register, the corresponding data address bit must match the value of the same bit position in the DAC1 register. For every bit set to 0 in the DAC2 register, the corresponding address bit comparison does not affect the result of the DAC event determination.

This comparison mode is useful for detecting accesses to a particular byte address, when the accesses may be of various sizes. For example, if the debugger is interested in detecting accesses to byte address 0x00000003, then these accesses may occur due to a byte access to that specific address, or due to a halfword access to address 0x00000002, or due to a word access to address 0x00000000. By using address bit mask mode and specifying that the low-order two bits of the address should be ignored (that is, setting the address bit mask in DAC2 to 0xFFFFF0), the debugger can detect each of these types of access to byte address 0x00000003.

When the data address matches the address bit mask mode conditions, either one or both of the DAC debug event bits corresponding to the operation type (read or write) will be set in the DBSR, as determined by which of the corresponding two DAC event enable bits are set in DBCR0. That is, when an address bit mask mode DAC debug event occurs, the setting of DBCR2[**DAC1R**, **DAC1W**, **DAC2R**, **DAC2W**] will determine whether one or the other or both of the DBSR[**DAC1R**, **DAC1W**, **DAC2R**, **DAC2W**] bits corresponding to the operation type are set. It is a programming error to set the DAC mode field to address bit mask mode without also enabling at least one of the four DAC event enable bits in DBCR0.

- Range inclusive comparison mode (DBCR2[**DAC12M**] = 0b10)

In this mode, the DAC1 or DAC2 event occurs only if the data address is *within* the range defined by the DAC1 and DAC2 register values, as follows:  $\text{DAC1} \leq \text{address} < \text{DAC2}$ .

When the data address falls within the specified range, either one or both of the DAC debug event bits corresponding to the operation type (read or write) will be set in the DBSR, as determined by which of the corresponding two DAC event enable bits are set in DBCR0. That is, when a range inclusive mode DAC debug event occurs, the setting of DBCR2[**DAC1R**, **DAC1W**, **DAC2R**, **DAC2W**] will determine whether one or the other or both of the DBSR[**DAC1R**, **DAC1W**, **DAC2R**, **DAC2W**] bits corresponding to the operation type are set. It is a programming error to set the DAC mode field to a range comparison mode (either inclusive or exclusive) without also enabling at least one of the four DAC event enable bits in DBCR0.

- Range exclusive comparison mode (DBCR2[**DAC12M**] = 0b11)

In this mode, the DAC1 or DAC2 event occurs only if the data address is *outside* the range defined by the DAC1 and DAC2 register values, as follows:  $\text{address} < \text{DAC1}$  or  $\text{address} \geq \text{DAC2}$ .

When the data address falls outside the specified range, either one or both of the DAC debug event bits corresponding to the operation type (read or write) will be set in the DBSR, as determined by which of the corresponding two DAC event enable bits are set in DBCR0. That is, when a range exclusive mode DAC debug event occurs, the setting of DBCR2[DAC1R, DAC1W, DAC2R, DAC2W] will determine whether one or the other or both of the DBSR[DAC1R, DAC1W, DAC2R, DAC2W] bits corresponding to the operation type are set. It is a programming error to set the DAC mode field to a range comparison mode (either inclusive or exclusive) without also enabling at least one of the four DAC event enable bits in DBCR0.

### DAC User/Supervisor Field

DBCR2[DAC1US, DAC2US] are the individual DAC user/supervisor fields for the two DAC events. The DAC user/supervisor fields specify what operating mode the processor must be in order for the corresponding DAC event to occur. The operating mode is determined by the Problem State field of the Machine State Register (MSR[PR]; see *User and Supervisor Modes* on page 78). When the DAC user/supervisor field is 0b00, the operating mode does not matter—the DAC debug event may occur independent of the state of MSR[PR]. When this field is 0b10, the processor must be operating in supervisor mode (MSR[PR] = 0). When this field is 0b11, the processor must be operating in user mode (MSR[PR] = 1). The DAC user/supervisor field value of 0b01 is reserved.

If the DAC mode is set to one of the “paired” modes (address bit mask mode, or one of the two range modes), it is a programming error (and the results of any data address comparison are undefined) if DBCR2[DAC1US] and DBCR2[DAC2US] are not set to the same value.

### DAC Effective/Real Address Field

DBCR2[DAC1ER, DAC2ER] are the individual DAC effective/real address fields for the two DAC events. The DAC effective/real address fields specify whether the instruction address comparison should be performed using the effective, virtual, or real address (see *Memory Management* on page 129) for an explanation of these different types of addresses). When the DAC effective/real address field is 0b00, the comparison is performed using the effective address only; the DAC debug event may occur independent of the data address space (MSR[DS]). When this field is 0b10, the DAC debug event occurs only if the effective address matches the DAC conditions and is in virtual address space 0 (MSR[DS] = 0). Similarly, when this field is 0b11, the DAC debug event occurs only if the effective address matches the DAC conditions and is in virtual address space 1 (MSR[DS] = 1). Note that in these latter two modes, in which the virtual address space of the data is considered, it is not the entire virtual address which is considered. The Process ID, which forms the final part of the virtual address, is not considered. Finally, the DAC effective/real address field value of 0b01 is reserved, and corresponds to the PowerPC Book-E architected real address comparison mode, which is not supported by the PPC440x5.

If the DAC mode is set to one of the “paired” modes (address bit mask mode, or one of the two range modes), it is a programming error (and the results of any data address comparison are undefined) if DBCR2[DAC1ER] and DBCR2[DAC2ER] are not set to the same value.

### DVC Byte Enable Field

DBCR2[DVC1BE, DVC2BE] are the individual data *value* compare (DVC) byte enable fields for the two DVC events. These fields must be disabled (by being set to 4b0000) in order for the corresponding DAC debug event to be enabled. In other words, when any of the DVC byte enable field bits for a given DVC event are set to 1, the corresponding DAC event is disabled, and the various DAC field conditions are used in conjunction with the DVC field conditions to determine whether a DVC event should occur. See *Data Value Compare (DVC) Debug Event* on page 225 for more information on DVC events.



### 8.3.2.2 DAC Debug Event Processing

The behavior of the PPC440x5 upon a DAC debug event depends on the setting of DBCR2[*DAC12A*]. This field of DBCR2 controls whether DAC debug events are processed in a *synchronous* (DBCR2[*DAC12A*] = 0) or an *asynchronous* (DBCR2[*DAC12A*] = 1) fashion.

#### DBCR2[*DAC12A*] = 0 (Synchronous Mode)

When operating in external debug mode or debug wait mode, the occurrence of a DAC debug event is recorded in the corresponding bit of the DBSR and causes the instruction execution to be suppressed. The processor then enters the stop state and ceases the processing of instructions. The program counter will contain the address of the instruction which caused the DAC debug event. Similarly, when operating in internal debug mode with Debug interrupts enabled (MSR[*DE*] = 1), the occurrence of a DAC debug event is recorded in the DBSR and causes the instruction execution to be suppressed. A Debug interrupt will occur with CSRR0 set to the address of the instruction which caused the DAC debug event.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled (MSR[*DE*] = 0), then a DAC debug event will set the corresponding DAC field of the DBSR, along with the Imprecise Debug Event (*IDE*) field of the DBSR. The instruction execution is not suppressed, as no Debug interrupt will occur immediately. Instead, instruction execution continues, and a Debug interrupt will occur if and when MSR[*DE*] is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the DAC debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the DBSR[*IDE*] field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of a DAC debug event simply sets the corresponding DAC field of the DBSR and is indicated over the trace interface, and instruction execution continues. DBCR2[*DAC12A*] does not affect the processing of DAC debug events when operating in trace mode.

**Engineering Note:** When DAC debug events are enabled in any debug mode other than trace mode, and DBCR2[*DAC12A*] is set to 0 (synchronous mode), in order for the PPC440x5 to deal with a DAC-related Debug interrupt in a synchronous fashion, the processing of all potential DAC debug event-causing instructions (loads, stores, and cache management instructions) is impacted by one processor cycle. This one cycle impact occurs whether or not the instruction is actually causing a DAC debug event. Overall processor performance is thus significantly impacted if synchronous mode DAC debug events are enabled. In order to maintain normal processor performance while DAC debug events are enabled *and in the absence of any actual DAC debug events*, software should set DBCR2[*DAC12A*] to 1.

#### DBCR2[*DAC12A*] = 1 (Asynchronous Mode)

When operating in external debug mode or debug wait mode, the occurrence of a DAC debug event is recorded in the corresponding bit of the DBSR and causes the processor to enter stop state and cease processing instructions. However, the determination and processing of the DAC debug event is not handled synchronously with respect to the instruction execution. That is, the processor may process the DAC debug event and enter the stop state either before or after the completion of the instruction causing the event. If the DAC debug event is processed *before* the completion of the instruction causing the event, then upon entering the stop state the program counter will contain the address of that instruction, and that

instruction's execution will have been suppressed. Conversely, if the DAC debug event is processed *after* the completion of the instruction causing the event, then the program counter will contain the address of some instruction after the one which caused the event. Whether or not the DAC debug event processing occurs before or after the completion of the instruction depends on the particular circumstances surrounding the instruction's execution, the details of which are generally beyond the scope of this document.

Similarly, when operating in internal debug mode with Debug interrupts enabled ( $MSR[DE] = 1$ ), the occurrence of a DAC debug event is recorded in the DBSR and will generate a Debug interrupt with  $CSRRO$  set to the address of the instruction which caused the DAC debug event, or to the address of some subsequent instruction, depending upon whether the event is processed before or after the instruction completes.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled ( $MSR[DE] = 0$ ), then a DAC debug event will set the corresponding DAC field of the DBSR, along with the Imprecise Debug Event (IDE) field of the DBSR. Instruction execution continues, and a Debug interrupt will occur if and when  $MSR[DE]$  is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the DAC debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the  $DBSR[IDE]$  field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of a DAC debug event simply sets the corresponding DAC field of the DBSR and is indicated over the trace interface, and instruction execution continues.  $DBCR2[DAC12A]$  does not affect the processing of DAC debug events when operating in trace mode.

### 8.3.2.3 DAC Debug Events Applied to Instructions that Result in Multiple Storage Accesses

Certain misaligned load and store instructions are handled by making multiple, independent storage accesses. Similarly, load and store multiple and string instructions which access more than one register result in more than one storage access. *Load and Store Alignment* on page 114 provides a detailed description of the circumstances that lead to such multiple storage accesses being made as the result of the execution of a single instruction.

Whenever the execution of a given instruction results in multiple storage accesses, the data address of each access is independently considered for whether or not it will cause a DAC debug event.

### 8.3.2.4 DAC Debug Events Applied to Various Instruction Types

Various special cases apply to the cache management instructions, the store word conditional indexed (**stwcx.**) instruction, and the load and store string indexed (**lswx**, **stswx**) instructions, with regards to DAC debug events. These special cases are as follows:

#### **dcbz, dcbi**

The **dcbz** and **dcbi** instructions are considered "stores" with respect to both storage access control and DAC debug events. The **dcbz** instruction directly changes the contents of a given storage location, whereas the **dcbi** instruction can indirectly change the contents of a given storage location by invalidating data which has been modified within the data cache, thereby "restoring" the value of the location to the "old" contents of memory. As "store" operations, they may cause DAC write debug events.



**dcbst, dcbf**

The **dcbst** and **dcbf** instructions are considered “loads” with respect to storage access control, since they do not change the contents of a given storage location. They may merely cause the data at that storage location to be moved from the data cache out to memory. However, in a debug environment, the fact that these instructions may lead to write operations on the external interface is typically the event of interest. Therefore, these instructions are considered “stores” with respect to DAC debug events, and may cause DAC write debug events.

**dcbt, dcbtst, icbt**

The touch instructions are considered “loads” with respect to both storage access control and DAC debug events. However, these instructions are treated as no-ops if they reference caching inhibited storage locations, or if they cause Data Storage or Data TLB Miss exceptions. Consequently, if a touch instruction is being treated as a no-op for one of these reasons, then it does not cause a DAC read debug event. However, if a touch instruction is not being treated as a no-op for one of these reasons, it may cause a DAC read debug event.

**dcba**

The **dcba** instruction is treated as a no-op by the PPC440x5, and thus will not cause a DAC debug event.

**icbi**

The **icbi** instruction is considered a “load” with respect to both storage access control and DAC debug events, and thus may cause a DAC read debug event.

**dccci, dcread, iccci, icread**

The **dccci** and **iccci** instructions do not generate an address, but rather they affect the entire data and instruction cache, respectively. Similarly, the **dcread** and **icread** instructions do not generate an address, but rather an “index” which is used to select a particular location in the respective cache, without regard to the storage address represented by that location. Therefore, none of these instructions cause DAC debug events.

**stwcx.**

If the execution of a **stwcx.** instruction would otherwise have caused a DAC write debug event, but the processor does not have the reservation from a **lwarx** instruction, then the DAC write debug event does not occur since the storage location does not get written.

**lswx, stswx**

DAC debug events do not occur for **lswx** or **stswx** instructions with a length of 0 ( $XER[TBC] = 0$ ), since these instructions do not actually access storage.

**8.3.3 Data Value Compare (DVC) Debug Event**

DVC debug events occur when execution is attempted of a load, store, or **dcbz** instruction for which the data storage address and other parameters match the DAC conditions specified by DBCR0, DBCR2, and the DAC registers, and for which the data accessed matches the DVC conditions specified by DBCR2 and the DVC

registers. In other words, in order for a DVC debug event to occur, the conditions for a DAC debug event must first be met, and then the data must also match the DVC conditions. *Data Address Compare (DAC) Debug Event* on page 220 describes the DAC conditions. In addition to the DAC conditions, there are two DVC registers on the PPC440x5, DVC1 and DVC2. The DVC registers can be used to specify two independent, 4-byte data values, which are selectively compared against the data being accessed by a given load, store, or cache management instruction.

When a DVC event occurs, the corresponding DBSR[DAC1R, DAC1W, DAC2R, DAC2W] bit is set. These same DBSR bits are shared by DAC debug events.

### 8.3.3.1 DVC Debug Event Fields

In addition to the DAC debug event fields described in *DAC Debug Event Fields* on page 220, and the DVC registers themselves, there are two fields in DBCR2 which are used to specify the DVC conditions, as follows:

#### DVC Byte Enable Field

DBCR2[DVC1BE, DVC2BE] are the individual DVC byte enable fields for the two DVC events. When one or the other (or both) of these fields is disabled (by being set to 4b0000), the corresponding DVC debug event is disabled (the corresponding DAC debug event may still be enabled, as determined by the DAC debug event enable field of DBCR0). When either one or both of these fields is enabled (by being set to a non-zero value), then the corresponding DVC debug event is enabled.

Each bit of a given DVC byte enable field corresponds to a byte position within an aligned word of memory. For a given aligned word of memory, the byte offsets (or “byte lanes”) within that word are numbered 0, 1, 2, and 3, starting from the left-most (most significant) byte of the word. Accordingly, bits 0:3 of a given DVC byte enable field correspond to bytes 0:3 of an aligned word of memory being accessed.

For an access to “match” the DVC conditions for a given byte, the access must be actually transferring data on that given byte position *and* the data must match the corresponding byte value within the DVC register.

For each storage access, the DVC comparison is made against the bytes that are being accessed within the aligned word of memory containing the starting byte of the transfer. For example, consider a load word instruction with a starting data address of x01. The four bytes from memory are located at addresses 0x01–0x04, but the aligned word of memory containing the starting byte consists of addresses 0x00–0x03. Thus the only bytes being accessed within the aligned word of memory containing the starting byte are the bytes at addresses 0x01–0x03, and only these bytes are considered in the DVC comparison. The byte transferred from address 0x04 is not considered.

#### DVC Mode Field

DBCR2[DVC1M, DVC2M] are the individual DVC mode fields for the two DVC events. Each one of these fields specifies the particular data value comparison mode for the corresponding DVC debug event. There are three comparison modes supported by the PPC440x5:

- AND comparison mode (DBCR2[DVC1M, DVC2M] = 0b01)

In this mode, all data byte lanes enabled by a DVC byte enable field must be being accessed and must match the corresponding byte data value in the corresponding DVC1 or DVC2 register.

- OR comparison mode (DBCR2[DVC1M, DVC2M] = 0b10)

In this mode, at least one data byte lane that is enabled by a DVC byte enable field must be being accessed and must match the corresponding byte data value in the corresponding DVC1 or DVC2 register.

- AND-OR comparison mode (DBCR2[DVC1M, DVC2M] = 0b11)

In this mode, the four byte lanes of an aligned word are divided into two pairs, with byte lanes 0 and 1 being in one pair, and byte lanes 2 and 3 in the other pair. The DVC comparison mode for each pair of byte lanes operates in AND mode, and then the results of these two AND mode comparisons are ORed together to determine whether a DVC debug event occurs. In other words, a DVC debug event occurs if either one or both of the pairs of byte lanes satisfy the AND mode comparison requirements.

This mode may be used to cause a DVC debug event upon an access of a particular halfword data value in either of the two halfwords of a word in memory.

### 8.3.3.2 DVC Debug Event Processing

The behavior of the PPC440x5 upon a DVC debug event depends on the setting of DBCR2[*DAC12A*]. This field of DBCR2 controls whether DVC debug events are processed in a *synchronous* (DBCR2[*DAC12A*] = 0) or an *asynchronous* (DBCR2[*DAC12A*] = 1) fashion. The processing of DVC debug events is the same as it is for DAC debug events. See *DAC Debug Event Processing* on page 223 for more information.

### 8.3.3.3 DVC Debug Events Applied to Instructions that Result in Multiple Storage Accesses

Certain misaligned load and store instructions are handled by making multiple, independent storage accesses. Similarly, load and store multiple and string instructions which access more than one register result in more than one storage access. *Load and Store Alignment* on page 114 provides a detailed description of the circumstances that lead to such multiple storage accesses being made as the result of the execution of a single instruction.

Whenever the execution of a given instruction results in multiple storage accesses, the address and data of each access is independently considered for whether or not it will cause a DVC debug event.

### 8.3.3.4 DVC Debug Events Applied to Various Instruction Types

Various special cases apply to the cache management instructions, the store word conditional indexed (**stwcx.**) instruction, and the load and store string indexed (**lswx**, **stswx**) instructions, with regards to DVC debug events. These special cases are as follows:

#### **dcbz**

The **dcbz** instruction is the only cache management instruction which can cause a DVC debug event. **dcbz** is the only such instruction which actually writes new data to a storage location (in this case, an entire 32-byte data cache line is written to zeroes).

#### **stwcx.**

If the execution of a **stwcx.** instruction would otherwise have caused a DVC write debug event, but the processor does not have the reservation from a **lwarx** instruction, then the DVC write debug event does not occur since the storage location does not get written.

## lswx, stswx

DVC debug events do not occur for **lswx** or **stswx** instructions with a length of 0 ( $XER[TBC] = 0$ ), since these instructions do not actually access storage.

### 8.3.4 Branch Taken (BRT) Debug Event

BRT debug events occur when BRT debug events are enabled ( $DBCR0[BRT] = 1$ ) and execution is attempted of a branch instruction for which the branch condition(s) are satisfied, such that the instruction stream will be redirected to the target address of the branch.

When operating in external debug mode or debug wait mode, the occurrence of a BRT debug event is recorded in  $DBSR[BRT]$  and causes the instruction execution to be suppressed. The processor then enters the stop state and ceases the processing of instructions. The program counter will contain the address of the branch instruction which caused the BRT debug event. Similarly, when operating in internal debug mode with Debug interrupts enabled ( $MSR[DE] = 1$ ), the occurrence of a BRT debug event is recorded in  $DBSR[BRT]$  and causes the instruction execution to be suppressed. A Debug interrupt will occur with  $CSRR0$  set to the address of the branch instruction which caused the BRT debug event.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled ( $MSR[DE] = 0$ ), then BRT debug events cannot occur. Since taken branches are a very common operation and thus likely to be frequently executed within the critical class interrupt handlers (which typically have  $MSR[DE]$  set to 0), allowing BRT debug events under these conditions would lead to an undesirable number of delayed (and hence imprecise) Debug interrupts.

When operating in trace mode, the occurrence of a BRT debug event is simply recorded in  $DBSR[BRT]$  and is indicated over the trace interface, and instruction execution continues.

### 8.3.5 Trap (TRAP) Debug Event

TRAP debug events occur when TRAP debug events are enabled ( $DBCR0[TRAP] = 1$ ) and execution is attempted of a trap (**tw**, **twi**) instruction for which the trap condition is satisfied.

When operating in external debug mode or debug wait mode, the occurrence of a TRAP debug event is recorded in  $DBSR[TRAP]$  and causes the instruction execution to be suppressed. The processor then enters the stop state and ceases the processing of instructions. The program counter will contain the address of the trap instruction which caused the TRAP debug event. Similarly, when operating in internal debug mode with Debug interrupts enabled ( $MSR[DE] = 1$ ), the occurrence of a TRAP debug event is recorded in  $DBSR[TRAP]$  and causes the instruction execution to be suppressed. A Debug interrupt will occur with  $CSRR0$  set to the address of the trap instruction which caused the TRAP debug event.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled ( $MSR[DE] = 0$ ), the occurrence of a TRAP debug event will set  $DBSR[TRAP]$ , along with the Imprecise Debug Event (IDE) field of the DBSR. Although a Debug interrupt will not occur immediately, the instruction execution is suppressed as a Trap exception type Program interrupt will occur instead. A Debug interrupt will also occur later, if and when  $MSR[DE]$  is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the TRAP debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the  $DBSR[IDE]$  field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of a TRAP debug event is simply recorded in  $DBSR[TRAP]$  and is indicated over the trace interface, and instruction execution continues.

### 8.3.6 Return (RET) Debug Event

RET debug events occur when RET debug events are enabled ( $DBCR0[RET] = 1$ ) and execution is attempted of a return (**rfi**, **rfci**, or **rfmci**) instruction.

When operating in external debug mode or debug wait mode, the occurrence of a RET debug event is recorded in  $DBSR[RET]$  and causes the instruction execution to be suppressed. The processor then enters the stop state and ceases the processing of instructions. The program counter will contain the address of the return instruction which caused the RET debug event. Similarly, when operating in internal debug mode with Debug interrupts enabled ( $MSR[DE] = 1$ ), the occurrence of a RET debug event is recorded in  $DBSR[RET]$  and causes the instruction execution to be suppressed. A Debug interrupt will occur with  $CSRR0$  set to the address of the return instruction which caused the RET debug event.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled ( $MSR[DE] = 0$ ), then RET debug events can occur only for **rfi** instructions, and not for **rfci** or **rfmci** instructions. Since the **rfci** or **rfmci** instruction is typically used to return from a critical class interrupt handler (including the Debug interrupt itself), and  $MSR[DE]$  is typically 0 at the time of the return, the **rfci** or **rfmci** must not be allowed to cause a RET debug event under these conditions, or else it would not be possible to return from the critical class interrupts.

For the **rfi** instruction only, if a RET debug event occurs under these conditions (internal debug mode enabled, external debug mode and debug wait mode disabled, and  $MSR[DE] = 0$ ), then  $DBSR[RET]$  is set, along with the Imprecise Debug Event (IDE) field of the DBSR. The instruction execution is not suppressed, as no Debug interrupt will occur immediately. Instead, instruction execution continues, and a Debug interrupt will occur if and when  $MSR[DE]$  is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the RET debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the  $DBSR[IDE]$  field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of a RET debug event is simply recorded in  $DBSR[RET]$  and is indicated over the trace interface, and instruction execution continues.

### 8.3.7 Instruction Complete (ICMP) Debug Event

ICMP debug events occur when ICMP debug events are enabled ( $DBCR0[ICMP] = 1$ ) and the PPC440x5 completes the execution of any instruction.

When operating in external debug mode or debug wait mode, the occurrence of an ICMP debug event is recorded in  $DBSR[ICMP]$  and causes the processor to enter the stop state and cease processing instructions. The program counter will contain the address of the instruction which would have executed next, had the ICMP debug event not occurred. Note that if the instruction whose completion caused the ICMP debug event was a branch instruction (and the branch conditions were satisfied), then upon entering the stop state the program counter will contain the target of the branch, and not the address of the instruction that is sequentially after the branch. Similarly, if the ICMP debug event is caused by the execution of a return (**rfi**, **rfci**, or **rfmci**) instruction, then upon entering the stop state the program counter will contain the address being *returned to*, and not the address of the instruction which is sequentially after the return instruction.

When operating in internal debug mode with Debug interrupts enabled ( $MSR[DE] = 1$ ), the occurrence of an ICMP debug event is recorded in  $DBSR[ICMP]$  and a Debug interrupt will occur with  $CSRR0$  set to the address of the instruction which would have executed next, had the ICMP debug event not occurred. Note

that there is a special case of  $MSR[DE] = 1$  at the time of the execution of the instruction causing the ICMP debug event, but that instruction itself sets  $MSR[DE]$  to 0. This special case is described in more detail in *Debug Interrupt* on page 188, in the subsection on the setting of CSRR0.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled ( $MSR[DE] = 0$ ), then ICMP debug events cannot occur. Since the code at the beginning of the critical class interrupt handlers (including the Debug interrupt itself) must execute at least temporarily with  $MSR[DE] = 0$ , there would be no way to avoid causing additional ICMP debug events and setting  $DBSR[IDE]$ , if ICMP debug events were allowed to occur under these conditions.

The PPC440x5 does not support the use of the ICMP debug event when operating in trace mode. Software must not enable ICMP debug events unless one of the other debug modes is enabled as well.

### 8.3.8 Interrupt (IRPT) Debug Event

IRPT debug events occur when IRPT debug events are enabled ( $DBCR0[IRPT] = 1$ ) and an interrupt occurs.

When operating in external debug mode or debug wait mode, the occurrence of an IRPT debug event is recorded in  $DBSR[IRPT]$  and causes the processor to enter the stop state and cease processing instructions. The program counter will contain the address of the instruction which would have executed next, had the IRPT debug event not occurred. Since the IRPT debug event is caused by the occurrence of an interrupt, by definition this address is that of the first instruction of the interrupt handler for the interrupt type which caused the IRPT debug event.

When operating in internal debug mode with external debug mode and debug wait mode both disabled (and regardless of the value of  $MSR[DE]$ ), an IRPT debug event can only occur due to a non-critical class interrupt. Critical class interrupts (Machine Check, Critical Input, Watchdog Timer, and Debug interrupts) cannot cause IRPT debug events in internal debug mode (unless also in external debug mode or debug wait mode), as otherwise the Debug interrupt which would occur as the result of the IRPT debug event would by necessity always be imprecise, since the critical class interrupt which would be causing the IRPT debug event would itself be causing  $MSR[DE]$  to be set to 0.

For a non-critical class interrupt which is causing an IRPT debug event while internal debug mode is enabled and external debug mode and debug wait mode are both disabled, the occurrence of the IRPT debug event is recorded in  $DBSR[IRPT]$ . If  $MSR[DE]$  is 1 at the time of the IRPT debug event, then a Debug interrupt occurs with CSRR0 set to the address of the instruction which would have executed next, had the IRPT debug event not occurred. Since the IRPT debug event is caused by the occurrence of some other interrupt, by definition this address is that of the first instruction of the interrupt handler for the interrupt type which caused the IRPT debug event. If  $MSR[DE]$  is 0 at the time of the IRPT debug event, then the Imprecise Debug Event (IDE) field of the DBSR is also set and a Debug interrupt does not occur immediately. Instead, instruction execution continues, and a Debug interrupt will occur if and when  $MSR[DE]$  is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the IRPT debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the  $DBSR[IDE]$  field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of an IRPT debug event is simply recorded in  $DBSR[IRPT]$  and is indicated over the trace interface, and instruction execution continues.



### 8.3.9 Unconditional Debug Event (UDE)

UDE debug events occur when a debug tool asserts the unconditional debug event request via the JTAG interface. The UDE debug event is the only event which does not have a corresponding enable field in DBCR0.

When operating in external debug mode or debug wait mode, the occurrence of a UDE debug event is recorded in DBSR[UDE] and causes the processor to enter the stop state and cease processing instructions. The program counter will contain the address of the instruction which would have executed next, had the UDE debug event not occurred. Similarly, when operating in internal debug mode with Debug interrupts enabled (MSR[DE] = 1), the occurrence of a UDE debug event is recorded in DBSR[UDE] and a Debug interrupt will occur with CSRR0 set to the address of the instruction which would have executed next, had the UDE debug event not occurred.

When operating in internal debug mode (and not also in external debug mode nor debug wait mode) with Debug interrupts disabled (MSR[DE] = 0), the occurrence of a UDE debug event will set DBSR[UDE], along with the Imprecise Debug Event (IDE) field of the DBSR. The Debug interrupt will not occur immediately. Instead, instruction execution continues, and a Debug interrupt will occur if and when MSR[DE] is set to 1, thereby enabling Debug interrupts, assuming software has not cleared the UDE debug event status from the DBSR in the meantime. Upon such a "delayed" interrupt, the Debug interrupt handler software may query the DBSR[IDE] field to determine that the Debug interrupt has occurred imprecisely.

When operating in trace mode, the occurrence of a UDE debug event simply sets DBSR[UDE] and is indicated over the trace interface, and instruction execution continues.

### 8.3.10 Debug Event Summary

Table 8-3 summarizes each of the debug event types, and the effect of debug mode and MSR[DE] on their occurrence.

Table 8-3. Debug Event Summary

External Debug Mode	Debug Wait Mode	Internal Debug Mode	MSR DE	Debug Events								
				IAC	DAC	DVC	BRT	TRAP	RET	ICMP	IRPT	UDE
Enabled	—	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
—	Enabled	—	—	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Disabled	Disabled	Enabled	1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Note 1	Yes

Table 8-3. Debug Event Summary (continued)

External Debug Mode	Debug Wait Mode	Internal Debug Mode	MSR DE	Debug Events								
				IAC	DAC	DVC	BRT	TRAP	RET	ICMP	IRPT	UDE
Disabled	Disabled	Enabled	0	Note 2	Yes	Yes	No	Yes	Note 3	No	Note 1	Yes
Disabled	Disabled	Disabled	—	Yes	Yes	Yes	Yes	Yes	Yes	Note 4	yes	Yes

## Table Notes

1. IRPT debug events may only occur for non-critical class interrupts when operating in internal debug mode with external debug mode and debug wait mode both disabled.
2. IAC debug events may not occur in internal debug mode with MSR[DE] = 0 and with external debug mode and debug wait mode both disabled, and the IAC mode set to range inclusive or range exclusive. They *may* occur if the IAC mode is set to exact.
3. RET debug events may not occur for **r<sub>fci</sub>** or **r<sub>fmci</sub>** instructions when operating in internal debug mode with MSR[DE] = 0 and with external debug mode and debug wait mode both disabled. They may only occur in this mode for the **r<sub>fi</sub>** instruction.
4. ICMP debug events are not permitted when operating in trace debug mode. Software must not enable ICMP debug events unless one of the other debug modes is enabled.

## 8.4 Debug Reset

Software can initiate an immediate reset operation by setting DBCR0[RST] to a non-zero value. The results of a reset operation within the PPC440x5 core are described in *Initialization* on page 83. The results of a reset operation on the rest of the chip and/or system is dependent on the particular type of reset operation (core, chip, or system reset), and on the particular chip and system implementation. See the chip user's manual for details.

## 8.5 Debug Timer Freeze

In order to maintain the semblance of “real time” operation while a system is being debugged, DBCR0[FT] can be set to 1, which will cause all of the timers within the PPC440x5 core to stop incrementing or decrementing for as long as a debug event bit is set in the DBSR, or until DBCR0[FT] is set to 0. See *Timer Facilities* on page 203 for more information on the operation of the PPC440x5 core timers.

## 8.6 Debug Registers

Various Special Purpose Registers (SPRs) are used to enable the debug modes, to configure and record debug events, and to communicate with debug tool hardware and software. These debug registers may be accessed either through software running on the processor or through the JTAG debug port of the PPC440x5.

**Programming Note:** It is the responsibility of software to synchronize the context of any changes to the debug facility registers. Specifically, when changing the contents of any of the debug facility registers, software must execute an **isync** instruction both *before* and *after* the changes to these registers, to ensure that all preceding instructions use the *old* values of the registers, and that all succeeding instructions use the *new* values. In addition, when



changing any of the debug facility register fields related to the DAC and/or DVC debug events, software must execute an **msync** instruction *before* making the changes, to ensure that all storage accesses complete using the *old* context of these register fields.

### 8.6.1 Debug Control Register 0 (DBCR0)

DBCR0 is an SPR that is used to enable debug modes and events, reset the processor, and control timer operation when debugging. DBCR0 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 8-1. Debug Control Register 0 (DBCR0)

0	EDM	External Debug Mode 0 Disable external debug mode. 1 Enable external debug mode.	
1	IDM	Internal Debug Mode 0 Disable internal debug mode. 1 Enable internal debug mode.	
2:3	RST	Reset 00 No action 01 Core reset 10 Chip reset 11 System reset  <b>Attention:</b> Writing 01, 10, or 11 to this field causes a processor reset to occur.	
4	ICMP	Instruction Completion Debug Event 0 Disable instruction completion debug event. 1 Enable instruction completion debug event.	Instruction completions do not cause instruction completion debug events if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
5	BRT	Branch Taken Debug Event 0 Disable branch taken debug event. 1 Enable branch taken debug event.	Taken branches do not cause branch taken debug events if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
6	IRPT	Interrupt Debug Event 0 Disable interrupt debug event. 1 Enable interrupt debug event.	Critical interrupts do not cause interrupt debug events in internal debug mode, unless also in external debug mode or debug wait mode.
7	TRAP	Trap Debug Event 0 Disable trap debug event. 1 Enable trap debug event.	
8	IAC1	Instruction Address Compare (IAC) 1 Debug Event 0 Disable IAC 1 debug event. 1 Enable IAC 1 debug event.	

9	IAC2	IAC 2 Debug Event 0 Disable IAC 2 debug event. 1 Enable IAC 2 debug event.	
10	IAC3	IAC 3 Debug Event 0 Disable IAC 3 debug event. 1 Enable IAC 3 debug event.	
11	IAC4	IAC 4 Debug Event 0 Disable IAC 4 debug event. 1 Enable IAC 4 debug event.	
12	DAC1R	Data Address Compare (DAC) 1 Read Debug Event 0 Disable DAC 1 read debug event. 1 Enable DAC 1 read debug event.	
13	DAC1W	DAC 1 Write Debug Event 0 Disable DAC 1 write debug event. 1 Enable DAC 1 write debug event.	
14	DAC2R	DAC 2 Read Debug Event 0 Disable DAC 2 read debug event. 1 Enable DAC 2 read debug event.	
15	DAC2W	DAC 2 Write Debug Event 0 Disable DAC 2 write debug event. 1 Enable DAC 2 write debug event.	
16	RET	Return Debug Event 0 Disable return ( <b>rfi/rfci/rfmc</b> i) debug event. 1 Enable return ( <b>rfi/rfci/rfmc</b> i) debug event.	<b>rfci/rfmc</b> i does not cause a return debug event if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
17:30		Reserved	
31	FT	Freeze timers on debug event 0 Timers are not frozen. 1 Freeze timers if a DBSR field associated with a debug event is set.	

### 8.6.2 Debug Control Register 1 (DBCR1)

DBCR1 is an SPR that is used to configure IAC debug events. DBCR1 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspr**.




Figure 8-2. Debug Control Register 1 (DBCR1)

0:1	IAC1US	Instruction Address Compare (IAC) 1 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)
-----	--------	---

2:3	IAC1ER	IAC 1 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
4:5	IAC2US	IAC 2 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
6:7	IAC2ER	IAC 2 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
8:9	IAC12M	IAC 1/2 Mode 00 Exact match 01 Reserved 10 Range inclusive 11 Range exclusive	Match if address[0:29] = IAC 1/2[0:29]; two independent compares  Match if $IAC1 \leq \text{address} < IAC2$ Match if $\text{address} < IAC1$ OR $\text{address} \geq IAC2$
10:14		Reserved	
15	IAC12AT	IAC 1/2 Auto-Toggle Enable 0 Disable IAC 1/2 auto-toggle 1 Enable IAC 1/2 auto-toggle	
16:17	IAC3US	IAC 3 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
18:19	IAC3ER	IAC 3 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
20:21	IAC4US	IAC 4 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
22:23	IAC4ER	IAC 4 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
24:25	IAC34M	IAC 3/4 Mode 00 Exact match 01 Reserved 10 Range inclusive 11 Range exclusive	Match if address[0:29] = IAC 3/4[0:29]; two independent compares  Match if $IAC3 \leq \text{address} < IAC4$ Match if $\text{address} < IAC3$ OR $\text{address} \geq IAC4$
26:30		Reserved	

---

31	IAC34AT	IAC3/4 Auto-Toggle Enable 0 Disable IAC 3/4 auto-toggle 1 Enable IAC 3/4 auto-toggle
----	---------	--

### 8.6.3 Debug Control Register 2 (DBCR2)

DBCR2 is an SPR that is used to configure DAC and DVC debug events. DBCR2 can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspr**.




Figure 8-3. Debug Control Register 2 (DBCR2)

0:1	DAC1US	Data Address Compare (DAC) 1 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
2:3	DAC1ER	DAC 1 Effective/Real 00 Effective (MSR[DS] = don't care) 01 Reserved 10 Virtual (MSR[DS] = 0) 11 Virtual (MSR[DS] = 1)	
4:5	DAC2US	DAC 2 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
6:7	DAC2ER	DAC 2 Effective/Real 00 Effective (MSR[DS] = don't care) 01 Reserved 10 Virtual (MSR[DS] = 0) 11 Virtual (MSR[DS] = 1)	
8:9	DAC12M	DAC 1/2 Mode 00 Exact match 01 Address bit mask 10 Range inclusive 11 Range exclusive	Match if address[0:31] = DAC 1/2[0:31]; two independent compares Match if address = DAC1; only compare bits corresponding to 1 bits in DAC2 Match if DAC1 ≤ address < DAC2 Match if address < DAC1 OR address ≥ DAC2
10	DAC12A	DAC 1/2 Asynchronous 0 Debug interrupt caused by DAC1/2 exception will be synchronous 1 Debug interrupt caused by DAC1/2 exception will be asynchronous	
11		Reserved	
12:13	DVC1M	Data Value Compare (DVC) 1 Mode 00 Reserved 01 AND all bytes enabled by DVC1BE 10 OR all bytes enabled by DVC1BE 11 AND-OR pairs of bytes enabled by DVC1BE	(0 AND 1) OR (2 AND 3)

14:15	DVC2M	DVC 2 Mode 00 Reserved 01 AND all bytes enabled by DVC2BE 10 OR all bytes enabled by DVC2BE 11 AND-OR pairs of bytes enabled by DVC2BE (0 AND 1) OR (2 AND 3)
16:19		Reserved
20:23	DVC1BE	DVC 1 Byte Enables 0:3
24:27		Reserved
28:31	DVC2BE	DVC 2 Byte Enables 0:3

### 8.6.4 Debug Status Register (DBSR)

The DBSR contains status on debug events as well as information on the type of the most recent reset. The status bits are set by the occurrence of debug events, while the reset type information is updated upon the occurrence of any of the three reset types.

The DBSR is read into a GPR using **mfspir**. Clearing the DBSR is performed using **mtspir** by placing a 1 in the GPR source register in all bit positions which are to be cleared in the DBSR, and a 0 in all other bit positions. The data written from the GPR to the DBSR is not direct data, but a mask. A 1 clears the bit and a 0 leaves the corresponding DBSR bit unchanged.




Figure 8-4. Debug Status Register (DBSR)

0	IDE	Imprecise Debug Event 0 Debug event occurred while MSR[DE] = 1 1 Debug event occurred while MSR[DE] = 0	For synchronous debug events in internal debug mode, this field indicates whether the corresponding Debug interrupt occurs precisely or imprecisely
1	UDE	Unconditional Debug Event 0 Event didn't occur 1 Event occurred	
2:3	MRR	Most Recent Reset 00 No reset has occurred since this field was last cleared by software. 01 Core reset 10 Chip reset 11 System reset	This field is set upon any processor reset to a value indicating the type of reset.
4	ICMP	Instruction Completion Debug Event 0 Event didn't occur 1 Event occurred	
5	BRT	Branch Taken Debug Event 0 Event didn't occur 1 Event occurred	

6	IRPT	Interrupt Debug Event 0 Event didn't occur 1 Event occurred
7	TRAP	Trap Debug Event 0 Event didn't occur 1 Event occurred
8	IAC1	IAC 1 Debug Event 0 Event didn't occur 1 Event occurred
9	IAC2	IAC 2 Debug Event 0 Event didn't occur 1 Event occurred
10	IAC3	IAC 3 Debug Event 0 Event didn't occur 1 Event occurred
11	IAC4	IAC 4 Debug Event 0 Event didn't occur 1 Event occurred
12	DAC1R	DAC 1 Read Debug Event 0 Event didn't occur 1 Event occurred
13	DAC1W	DAC 1 Write Debug Event 0 Event didn't occur 1 Event occurred
14	DAC2R	DAC 2 Read Debug Event 0 Event didn't occur 1 Event occurred
15	DAC2W	DAC 2 Write Debug Event 0 Event didn't occur 1 Event occurred
16	RET	Return Debug Event 0 Event didn't occur 1 Event occurred
17:29		Reserved
30	IAC12ATS	IAC 1/2 Auto-Toggle Status 0 Range is not reversed from value specified in DBCR1[IAC12M] 1 Range is reversed from value specified in DBCR1[IAC12M]
31	IAC34ATS	IAC 3/4 Auto-Toggle Status 0 Range is not reversed from value specified in DBCR1[IAC34M] 1 Range is reversed from value specified in DBCR1[IAC34M]

### 8.6.5 Instruction Address Compare Registers (IAC1–IAC4)

The four IAC registers specify the addresses upon which IAC debug events should occur. Each of the IAC registers can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 8-5. Instruction Address Compare Registers (IAC1–IAC4)

0:29		Instruction Address Compare (IAC) word address
30:31		Reserved

### 8.6.6 Data Address Compare Registers (DAC1–DAC2)

The two DAC registers specify the addresses upon which DAC (and/or DVC) debug events should occur. Each of the DAC registers can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.




Figure 8-6. Data Address Compare Registers (DAC1–DAC2)

0:31		Data Address Compare (DAC) byte address
------	--	---

### 8.6.7 Data Value Compare Registers (DVC1–DVC2)

The DVC registers specify the data values upon which DVC debug events should occur. Each of the DVC registers can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.





Figure 8-7. Data Value Compare Registers (DVC1–DVC2)

0:31		Data value to compare
------	--	-----------------------




### 8.6.8 Debug Data Register (DBDR)

The DBDR can be used for communication between software running on the processor and debug tool hardware and software. The DBDR can be written from a GPR using **mtspr**, and can be read into a GPR using **mfspir**.



*Figure 8-8. Debug Data Register (DBDR)*





## 9. Instruction Set

Descriptions of the PPC440x5 instructions follow. Each description contains the following elements:

- Instruction names (mnemonic and full)
- Instruction syntax
- Instruction format diagram
- Pseudocode description
- Prose description
- Registers altered

Where appropriate, instruction descriptions list invalid instruction forms and exceptions, and provide programming notes.

Table 9-1 summarizes the PPC440x5 instruction set by category.

Table 9-1. Instruction Categories

Category	Sub-Category	Instruction Types
<b>Integer</b>	Integer Storage Access	load, store
	Integer Arithmetic	add, subtract, multiply, divide, negate
	Integer Logical	and, andc, or, orc, xor, nand, nor, xnor, extend sign, count leading zeros
	Integer Compare	compare, compare logical
	Integer Trap	trap
	Integer Rotate	rotate and insert, rotate and mask
	Integer Shift	shift left, shift right, shift right algebraic
	Integer Select	select operand
<b>Branch</b>		branch, branch conditional, branch to link, branch to count
<b>Processor Control</b>	Condition Register Logical	crand, crandc, cror, crorc, crnand, crnor, crxor, crxnor
	Register Management	move to/from SPR, move to/from DCR, move to/from MSR, write to external interrupt enable bit, move to/from CR
	System Linkage	system call, return from interrupt, return from critical interrupt, return from machine check interrupt
	Processor Synchronization	instruction synchronize
<b>Storage Control</b>	Cache Management	data allocate, data invalidate, data touch, data zero, data flush, data store, instruction invalidate, instruction touch
	TLB Management	read, write, search, synchronize
	Storage Synchronization	memory synchronize, memory barrier
<b>Allocated</b>	Allocated Arithmetic	multiply-accumulate, negative multiply-accumulate, multiply halfword
	Allocated Logical	detect left-most zero byte
	Allocated Cache Management	data congruence-class invalidate, instruction congruence-class invalidate
	Allocated Cache Debug	data read, instruction read

## 9.1 Instruction Set Portability

To support embedded real-time applications, the PPC440x5 core implements the defined instruction set of the Book-E Enhanced PowerPC Architecture, with the exception of those operations which are defined for 64-bit implementations only, and those which are defined as floating-point operations. Support for the floating-point operations is provided via the auxiliary processor interface, while the 64-bit operations are not supported at all. See *Instruction Classes* on page 52 for more information on the support for defined instructions within the PPC440x5.

The PPC440x5 core also implements a number of instructions that are not part of PowerPC Book-E architecture, but are included as part of the PPC440x5. Architecturally, they are considered allocated instructions, as they use opcodes which are within the allocated class of instructions, which the PowerPC Book-E architecture identifies as being available for implementation-dependent and/or application-specific purposes. However, all of the allocated instructions which are implemented within the PPC440x5 core are “standard” for IBM PowerPC 400 Series family of embedded controllers, and are not unique to the PPC440x5.

The allocated instructions implemented within the PPC440x5 are divided into four sub-categories, and are shown in *Table 9-2*. Programs using these instructions may not be portable to other PowerPC Book-E implementations.

*Table 9-2. Allocated Instructions*

Arithmetic			Logical	Cache Management	Cache Debug
Multiply-Accumulate	Negative Multiply-Accumulate	Multiply Halfword			
<b>macchw[o][.]</b> <b>macchws[o][.]</b> <b>macchwsu[o][.]</b> <b>macchw[o][.]</b> <b>machhw[o][.]</b> <b>machhws[o][.]</b> <b>machhwsu[o][.]</b> <b>machhwu[o][.]</b> <b>maclhw[o][.]</b> <b>maclhws[o][.]</b> <b>maclhwsu[o][.]</b> <b>maclhwu[o][.]</b>	<b>nmacchw[o][.]</b> <b>nmacchws[o][.]</b> <b>nmachhw[o][.]</b> <b>nmachhws[o][.]</b> <b>nmaclhw[o][.]</b> <b>nmaclhws[o][.]</b>	<b>mulchw[.]</b> <b>mulchwu[.]</b> <b>mulhhw[.]</b> <b>mulhhwu[.]</b> <b>mullhw[.]</b> <b>mullhwu[.]</b>	<b>dImzb[.]</b>	<b>dccci</b> <b>iccci</b>	<b>dcread</b> <b>icread</b>

## 9.2 Instruction Formats

For more detailed information about instruction formats, including a summary of instruction field usage and instruction format diagrams for the PPC440x5 core, see *Instruction Formats* on page 507.

Instructions are four bytes long. Instruction addresses are always word-aligned.

Instruction bits 0 through 5 always contain the primary opcode. Many instructions have an extended opcode field as well. The remaining instruction bits contain additional fields. All instruction fields belong to one of the following categories:

- Defined

These instruction fields contain values, such as opcodes, that cannot be altered. The instruction format diagrams specify the values of defined fields.

- Variable

These fields contain operands, such as general purpose register specifiers and immediate values, each of which may contain any one of a number of values. The instruction format diagrams specify the field names of variable fields.

- Reserved

Bits in a reserved field should be set to 0. In the instruction format diagrams, reserved fields are shaded.

If any bit in a defined field does not contain the specified value, the instruction is illegal and an Illegal Instruction exception type Program interrupt occurs. If any bit in a reserved field does not contain 0, the instruction form is invalid and its result is architecturally undefined. Unless otherwise noted, the PPC440x5 core will execute all invalid instruction forms without causing an Illegal Instruction exception.

### 9.3 Pseudocode

The pseudocode that appears in the instruction descriptions provides a semi-formal language for describing instruction operations.

The pseudocode uses the following notation:

+	Twos complement addition
%	Remainder of an integer division; (33 % 32) = 1.
$\overset{u}{<}, \overset{u}{>}$	Unsigned comparison relations
(GPR(r))	The contents of GPR r, where $0 \leq r \leq 31$ .
(RA 0)	The contents of the register RA or 0, if the RA field is 0.
(Rx)	The contents of a GPR, where X is A, B, S, or T
0bn	A binary number
0xn	A hexadecimal number
<, >	Signed comparison relations
=	Assignment
=, ≠	Equal, not equal relations
CEIL(x)	Least integer $\geq x$ .
CIA	Current instruction address; the 32-bit address of the instruction being described by a sequence of pseudocode. This address is used to set the next instruction address (NIA). Does not correspond to any architected register.
DCR(DCRN)	A Device Control Register (DCR) specified by the DCRF field in an <b>mfocr</b> or <b>mtocr</b> instruction
EA	Effective address; the 32-bit address, derived by applying indexing or indirect addressing rules to the specified operand, that specifies an location in main storage.
EXTS(x)	The result of extending X on the left with sign bits.
FLD	An instruction or register field

FLD <sub>b</sub>	A bit in a named instruction or register field
FLD <sub>b,b,...</sub>	A list of bits, by number or name, in a named instruction or register field
FLD <sub>b:b</sub>	A range of bits in a named instruction or register field
GPR( <i>r</i> )	General Purpose Register (GPR) <i>r</i> , where $0 \leq r \leq 31$ .
GPRs	RA, RB, ...
MASK(MB,ME)	Mask having 1s in positions MB through ME (wrapping if MB > ME) and 0s elsewhere.
MS(addr, <i>n</i> )	The number of bytes represented by <i>n</i> at the location in main storage represented by <i>addr</i> .
NIA	Next instruction address; the 32-bit address of the next instruction to be executed. In pseudocode, a successful branch is indicated by assigning a value to NIA. For instructions that do not branch, the NIA is CIA +4.
PC	Program counter.
REG[FLD, FLD ...]	A list of fields in a named register
REG[FLD:FLD]	A range of fields in a named register
REG[FLD]	A field in a named register
REG <sub>b</sub>	A bit in a named register
REG <sub>b,b,...</sub>	A list of bits, by number or name, in a named register
REG <sub>b:b</sub>	A range of bits in a named register
RESERVE	Reserve bit; indicates whether a process has reserved a block of storage.
ROTL((RS), <i>n</i> )	Rotate left; the contents of RS are shifted left the number of bits specified by <i>n</i> .
SPR(SPRN)	A Special Purpose Register (SPR) specified by the SPRF field in an <b>mf spr</b> or <b>mt spr</b> instruction
c <sub>0:3</sub>	A four-bit object used to store condition results in compare instructions.
do	Do loop. "to" and "by" clauses specify incrementing an iteration variable; "while" and "until" clauses specify terminating conditions. Indenting indicates the scope of a loop.
if...then...else...	Conditional execution; if <i>condition</i> then <i>a</i> else <i>b</i> , where <i>a</i> and <i>b</i> represent one or more pseudocode statements. Indenting indicates the ranges of <i>a</i> and <i>b</i> . If <i>b</i> is null, the else does not appear.
instruction(EA)	An instruction operating on a data or instruction cache block associated with an EA.
leave	Leave innermost do loop or do loop specified in a leave statement.
<i>n</i>	A decimal number
<sup><i>n</i></sup> <i>b</i>	The bit or bit value <i>b</i> is replicated <i>n</i> times.
xx	Bit positions which are don't-cares.
	Concatenation
×	Multiplication
÷	Division yielding a quotient
⊕	Exclusive-OR (XOR) logical operator
–	Twos complement subtraction, unary minus
¬	NOT logical operator
^	AND logical operator
∨	OR logical operator

### 9.3.1 Operator Precedence

Table 9-3 lists the pseudocode operators and their associativity in descending order of precedence:

Table 9-3. Operator Precedence

Operators	Associativity
REG <sub>b</sub> , REG[FLD], function evaluation	Left to right
n <sub>b</sub>	Right to left
¬, − (unary minus)	Right to left
×, ÷	Left to right
+, −	Left to right
	Left to right
=, ≠, <, >, < <sup>u</sup> , > <sup>u</sup>	Left to right
∧, ⊕	Left to right
∨	Left to right
←	None

### 9.4 Register Usage

Each instruction description lists the registers altered by the instruction. Some register changes are explicitly detailed in the instruction description (for example, the target register of a load instruction). Some instructions also change other registers, but the details of the changes are not included in the instruction descriptions. Common examples of these kinds of register changes include the Condition Register (CR) and the Integer Exception Register (XER). For discussion of the CR, see *Condition Register (CR)* on page 66. For discussion of the XER, see *Integer Exception Register (XER)* on page 70.

### 9.5 Alphabetical Instruction Listing


The following pages list the instructions, both defined and allocated, which are implemented within the PPC440x5 core.

# add

Add



<b>add</b>	RT, RA, RB	OE=0, Rc=0
<b>add.</b>	RT, RA, RB	OE=0, Rc=1
<b>addo</b>	RT, RA, RB	OE=1, Rc=0
<b>addo.</b>	RT, RA, RB	OE=1, Rc=1



$$(RT) \leftarrow (RA) + (RB)$$

The sum of the contents of register RA and the contents of register RB is placed into register RT.


### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1





<b>addc</b>	RT, RA, RB	OE=0, Rc=0
<b>addc.</b>	RT, RA, RB	OE=0, Rc=1
<b>addco</b>	RT, RA, RB	OE=1, Rc=0
<b>addco.</b>	RT, RA, RB	OE=1, Rc=1



```

(RT) ← (RA) + (RB)
if (RA) + (RB)  $\overset{u}{>} 2^{32} - 1$  then
  XER[CA] ← 1
else
  XER[CA] ← 0

```

The sum of the contents of register RA and register RB is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.

**Registers Altered**

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

## adde


Add Extended

PPC440x5 CPU Core User's Manual



Preliminary

<b>adde</b>	RT, RA, RB	OE=0, Rc=0
<b>adde.</b>	RT, RA, RB	OE=0, Rc=1
<b>addeo</b>	RT, RA, RB	OE=1, Rc=0
<b>addeo.</b>	RT, RA, RB	OE=1, Rc=1




```
(RT) ← (RA) + (RB) + XER[CA]
if (RA) + (RB) + XER[CA]  $\geq 2^{32} - 1$  then
    XER[CA] ← 1
else
    XER[CA] ← 0
```

The sum of the contents of register RA, register RB, and XER[CA] is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.

### Registers Altered

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**addi** RT, RA, IM


$$(RT) \leftarrow (RA|0) + \text{EXTS}(IM)$$

If the RA field is 0, the IM field, sign-extended to 32 bits, is placed into register RT.

If the RA field is nonzero, the sum of the contents of register RA and the contents of the IM field, sign-extended to 32 bits, is placed into register RT.

### Registers Altered

- RT

### Programming Note

To place an immediate, sign-extended value into the GPR specified by RT, set RA = 0.

*Table 9-4. Extended Mnemonics for addi*

Mnemonic	Operands	Function	Other Registers Altered
<b>la</b>	RT, D(RA)	Load address (RA ≠ 0); D is an offset from a base address that is assumed to be (RA). $(RT) \leftarrow (RA) + \text{EXTS}(D)$ <i>Extended mnemonic for</i> <b>addi RT,RA,D</b>	
<b>li</b>	RT, IM	Load immediate. $(RT) \leftarrow \text{EXTS}(IM)$ <i>Extended mnemonic for</i> <b>addi RT,0,IM</b>	
<b>subi</b>	RT, RA, IM	Subtract EXTS(IM) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> <b>addi RT,RA,-IM</b>	


## addic

Add Immediate Carrying  
PPC440x5 CPU Core User's Manual



Preliminary

**addic** RT, RA, IM



```
(RT) ← (RA) + EXTS(IM)
if (RA) + EXTS(IM)  $\geq 2^{32} - 1$  then
    XER[CA] ← 1
else
    XER[CA] ← 0
```

The sum of the contents of register RA and the contents of the IM field, sign-extended to 32 bits, is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.


### Registers Altered

- RT
- XER[CA]

Table 9-5. Extended Mnemonics for addic

Mnemonic	Operands	Function	Other Registers Altered
<b>subic</b>	RT, RA, IM	Subtract EXTS(IM) from (RA) Place result in RT; place carry-out in XER[CA]. <i>Extended mnemonic for</i> <b>addic RT,RA,-IM</b>	

**addic.** RT, RA, IM



```
(RT) ← (RA) + EXTS(IM)
if (RA) + EXTS(IM) ≥ 232 - 1 then
    XER[CA] ← 1
else
    XER[CA] ← 0
```

The sum of the contents of register RA and the contents of the IM field, sign-extended to 32 bits, is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.

**Registers Altered**

- RT
- XER[CA]
- CR[CR0]

**Programming Note**

**addic.** is one of three instructions that implicitly update CR[CR0] without having an RC field. The other instructions are **andi.** and **andis.**

*Table 9-6. Extended Mnemonics for addic.*

Mnemonic	Operands	Function	Other Registers Altered
<b>subic.</b>	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT; place carry-out in XER[CA]. <i>Extended mnemonic for addic. RT,RA,-IM</i>	CR[CR0]

# addis


Add Immediate Shifted

PPC440x5 CPU Core User's Manual



Preliminary

**addis** RT, RA, IM



$$(RT) \leftarrow (RA|0) + (IM \parallel 16_0)$$

If the RA field is 0, the IM field is concatenated on its right with sixteen 0-bits and placed into register RT.

If the RA field is nonzero, the contents of register RA are added to the contents of the extended IM field. The sum is stored into register RT.

### Registers Altered

- RT

### Programming Note

An **addi** instruction stores a sign-extended 16-bit value in a GPR. An **addis** instruction followed by an **ori** instruction stores an arbitrary 32-bit value in a GPR, as shown in the following example:


```
addis    RT, 0, high 16 bits of value
ori      RT, RT, low 16 bits of value
```

Table 9-7. Extended Mnemonics for addis

Mnemonic	Operands	Function	Other Registers Altered
<b>lis</b>	RT, IM	Load immediate shifted. $(RT) \leftarrow (IM \parallel 16_0)$ <i>Extended mnemonic for</i> <b>addis RT,0,IM</b>	
<b>subis</b>	RT, RA, IM	Subtract $(IM \parallel 16_0)$ from $(RA 0)$ . Place result in RT. <i>Extended mnemonic for</i> <b>addis RT,RA,-IM</b>	

## Preliminary

<b>addme</b>	RT, RA	OE=0, Rc=0
<b>addme.</b>	RT, RA	OE=0, Rc=1
<b>addmeo</b>	RT, RA	OE=1, Rc=0
<b>addmeo.</b>	RT, RA	OE=1, Rc=1



```

(RT) ← (RA) + XER[CA] + (-1)
if (RA) + XER[CA] + 0xFFFF FFFF > 232 - 1 then
  XER[CA] ← 1
else
  XER[CA] ← 0
  
```

The sum of the contents of register RA, XER[CA], and -1 is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.

**Registers Altered**

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Invalid Instruction Forms**

- Reserved fields

## addze


Add to Zero Extended

PPC440x5 CPU Core User's Manual



Preliminary

<b>addze</b>	RT, RA	OE=0, Rc=0
<b>addze.</b>	RT, RA	OE=0, Rc=1
<b>addzeo</b>	RT, RA	OE=1, Rc=0
<b>addzeo.</b>	RT, RA	OE=1, Rc=1



```
(RT) ← (RA) + XER[CA]
if (RA) + XER[CA]  $\geq 2^{32} - 1$  then
  XER[CA] ← 1
else
  XER[CA] ← 0
```

The sum of the contents of register RA and XER[CA] is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the add operation.

### Registers Altered

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### Invalid Instruction Forms

- Reserved fields





**and** RA, RS, RB Rc=0  
**and.** RA, RS, RB Rc=1



$$(RA) \leftarrow (RS) \wedge (RB)$$

The contents of register RS are ANDed with the contents of register RB; the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if Rc contains 1

## andc


AND with Complement

PPC440x5 CPU Core User's Manual



Preliminary

**andc**            RA,RS,RB                            Rc=0  
**andc.**            RA,RS,RB                            Rc=1




$$(RA) \leftarrow (RS) \wedge \neg(RB)$$

The contents of register RS are ANDed with the ones complement of the contents of register RB; the result is placed into register RA.

### Registers Altered

- RA
- CR[CR0] if Rc contains 1

**andi.** RA, RS, IM



$$(RA) \leftarrow (RS) \wedge (^{16}0 \parallel IM)$$

The IM field is extended to 32 bits by concatenating 16 0-bits on its left. The contents of register RS is ANDed with the extended IM field; the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0]

**Programming Note**

The **andi.** instruction can test whether any of the 16 least-significant bits in a GPR are 1-bits.

**andi.** is one of three instructions that implicitly update CR[CR0] without having an Rc field. The other instructions are **addic.** and **andis..**


## andis.

AND Immediate Shifted  
PPC440x5 CPU Core User's Manual



Preliminary

**andis.** RA, RS, IM



$$(RA) \leftarrow (RS) \wedge (IM \parallel ^{16}0)$$

The IM field is extended to 32 bits by concatenating 16 0-bits on its right. The contents of register RS are ANDed with the extended IM field; the result is placed into register RA.

### Registers Altered

- RA
- CR[CR0]

### Programming Note

The **andis.** instruction can test whether any of the 16 most-significant bits in a GPR are 1-bits.

**andis.** is one of three instructions that implicitly update CR[CR0] without having an Rc field. The other instructions are **addic.** and **andi..**



<b>b</b>	target	AA=0, LK=0
<b>ba</b>	target	AA=1, LK=0
<b>bl</b>	target	AA=0, LK=1
<b>bla</b>	target	AA=1, LK=1

18	LI	AA	LK
0	6	30	31

```

If AA = 1 then
  LI ← target6:29
  NIA ← EXTS(LI || 20)
else
  LI ← (target - CIA)6:29
  NIA ← CIA + EXTS(LI || 20)
if LK = 1 then
  (LR) ← CIA + 4
PC ← NIA

```

The next instruction address (NIA) is the effective address of the branch target. The NIA is formed by adding a displacement to a base address. The displacement is obtained by concatenating two 0-bits to the right of the LI field and sign-extending the result to 32 bits.

If the AA field contains 0, the base address is the address of the branch instruction, which is the current instruction address (CIA). If the AA field contains 1, the base address is 0.

Instruction execution resumes with the instruction at the NIA.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

### Registers Altered

- LR if LK contains 1

## bc

Branch Conditional  
PPC440x5 CPU Core User's Manual



Preliminary

<b>bc</b>	BO, BI, target	AA=0, LK=0
<b>bca</b>	BO, BI, target	AA=1, LK=0
<b>bcl</b>	BO, BI, target	AA=0, LK=1
<b>bcla</b>	BO, BI, target	AA=1, LK=1

16		BO		BI		BD		AA	LK
0	6	11	16	30	31				

```

if BO2 = 0 then
    CTR ← CTR - 1
if (BO2 = 1 ∨ ((CTR = 0) = BO3)) ∧ (BO0 = 1 ∨ (CRBI = BO1)) then
    if AA = 1 then
        BD ← target16:29
        NIA ← EXTS(BD || 20)
    else
        BD ← (target - CIA)16:29
        NIA ← CIA + EXTS(BD || 20)
    else
        NIA ← CIA + 4
if LK = 1 then
    (LR) ← CIA + 4
PC ← NIA

```

If BO<sub>2</sub> contains 0, the CTR decrements, and the decremented value is tested for 0 as part of the branch condition. In this case, BO<sub>3</sub> indicates whether the test for 0 must be true or false in order for the branch to be taken. If BO<sub>2</sub> contains 1, then the CTR is neither decremented nor tested as part of the branch condition.

If BO<sub>0</sub> contains 0, then the CR bit specified by the BI field is compared to BO<sub>1</sub> as part of the branch condition. If BO<sub>0</sub> contains 1, then the CR is not tested as part of the branch condition, and the BI field is ignored.

The next instruction address (NIA) is either the effective address of the branch target, or the address of the instruction after the branch, depending on whether the branch is taken or not. The branch target address is formed by adding a displacement to a base address. The displacement is obtained by concatenating two 0-bits to the right of the BD field and sign-extending the result to 32 bits.

If the AA field contains 0, the base address is the address of the branch instruction, which is the current instruction address (CIA). If the AA field contains 1, the base address is 0.

BO<sub>4</sub> affects branch prediction, a performance-improvement feature. See *Branch Prediction* on page 64 for a complete discussion.

Instruction execution resumes with the instruction at the NIA.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

### Registers Altered

- CTR if BO<sub>2</sub> contains 0
- LR if LK contains 1

Table 9-8. Extended Mnemonics for bc, bca, bcl, bcla

Mnemonic	Operands	Function	Other Registers Altered
<b>bdnz</b>	target	Decrement CTR; branch if CTR $\neq$ 0. <i>Extended mnemonic for</i> <b>bc 16,0,target</b>	
<b>bdnza</b>		<i>Extended mnemonic for</i> <b>bca 16,0,target</b>	
<b>bdnzl</b>		<i>Extended mnemonic for</i> <b>bcl 16,0,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdnzla</b>		<i>Extended mnemonic for</i> <b>bcla 16,0,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdnzf</b>	cr_bit, target	Decrement CTR. Branch if CTR $\neq$ 0 AND CR <sub>cr_bit</sub> = 0. <i>Extended mnemonic for</i> <b>bc 0,cr_bit,target</b>	
<b>bdnzfa</b>		<i>Extended mnemonic for</i> <b>bca 0,cr_bit,target</b>	
<b>bdnzfl</b>		<i>Extended mnemonic for</i> <b>bcl 0,cr_bit,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdnzfla</b>		<i>Extended mnemonic for</i> <b>bcla 0,cr_bit,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdnztt</b>	cr_bit, target	Decrement CTR. Branch if CTR $\neq$ 0 AND CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 8,cr_bit,target</b>	
<b>bdnztta</b>		<i>Extended mnemonic for</i> <b>bca 8,cr_bit,target</b>	
<b>bdnzttl</b>		<i>Extended mnemonic for</i> <b>bcl 8,cr_bit,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdnzttla</b>		<i>Extended mnemonic for</i> <b>bcla 8,cr_bit,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdz</b>	target	Decrement CTR; branch if CTR = 0. <i>Extended mnemonic for</i> <b>bc 18,0,target</b>	
<b>bdza</b>		<i>Extended mnemonic for</i> <b>bca 18,0,target</b>	
<b>bdzl</b>		<i>Extended mnemonic for</i> <b>bcl 18,0,target</b>	(LR) $\leftarrow$ CIA + 4.
<b>bdzla</b>		<i>Extended mnemonic for</i> <b>bcla 18,0,target</b>	(LR) $\leftarrow$ CIA + 4.

Table 9-8. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bdzf</b>	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 0. <i>Extended mnemonic for</i> <b>bc 2,cr_bit,target</b>	
<b>bdzfa</b>		<i>Extended mnemonic for</i> <b>bca 2,cr_bit,target</b>	
<b>bdzfl</b>		<i>Extended mnemonic for</i> <b>bcl 2,cr_bit,target</b>	(LR) ← CIA + 4.
<b>bdzfla</b>		<i>Extended mnemonic for</i> <b>bcla 2,cr_bit,target</b>	(LR) ← CIA + 4.
<b>bdzt</b>	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 10,cr_bit,target</b>	
<b>bdzta</b>		<i>Extended mnemonic for</i> <b>bca 10,cr_bit,target</b>	
<b>bdztl</b>		<i>Extended mnemonic for</i> <b>bcl 10,cr_bit,target</b>	(LR) ← CIA + 4.
<b>bdztlia</b>		<i>Extended mnemonic for</i> <b>bcla 10,cr_bit,target</b>	(LR) ← CIA + 4.
<b>beq</b>	[cr_field,] target	Branch if equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+2,target</b>	
<b>beqa</b>		<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+2,target</b>	
<b>beql</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+2,target</b>	(LR) ← CIA + 4.
<b>beqla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+2,target</b>	(LR) ← CIA + 4.
<b>bf</b>	cr_bit, target	Branch if CR <sub>cr_bit</sub> = 0. <i>Extended mnemonic for</i> <b>bc 4,cr_bit,target</b>	
<b>bfa</b>		<i>Extended mnemonic for</i> <b>bca 4,cr_bit,target</b>	
<b>bfl</b>		<i>Extended mnemonic for</i> <b>bcl 4,cr_bit,target</b>	LR
<b>bfla</b>		<i>Extended mnemonic for</i> <b>bcla 4,cr_bit,target</b>	LR



Table 9-8. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bge</b>		Branch if greater than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+0,target</b>	
<b>bgea</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+0,target</b>	
<b>bgel</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+0,target</b>	LR
<b>bgea</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+0,target</b>	LR
<b>bgt</b>		Branch if greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+1,target</b>	
<b>bgt</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+1,target</b>	
<b>bgtl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+1,target</b>	LR
<b>bgtla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+1,target</b>	LR
<b>ble</b>		Branch if less than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+1,target</b>	
<b>blea</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+1,target</b>	
<b>blel</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+1,target</b>	LR
<b>blea</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+1,target</b>	LR
<b>blt</b>		Branch if less than Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+0,target</b>	
<b>blt</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+0,target</b>	
<b>bltl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+0,target</b>	(LR) ← CIA + 4.
<b>bltla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+0,target</b>	(LR) ← CIA + 4.

Table 9-8. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bne</b>		Branch if not equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+2,target</b>	
<b>bnea</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+2,target</b>	
<b>bnel</b>		<b>Extended mnemonic for</b> <b>bcl 4,4*cr_field+2,target</b>	(LR) ← CIA + 4.
<b>bnela</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+2,target</b>	(LR) ← CIA + 4.
<b>bng</b>		Branch if not greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+1,target</b>	
<b>bnga</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+1,target</b>	
<b>bngl</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+1,target</b>	(LR) ← CIA + 4.
<b>bngla</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+1,target</b>	(LR) ← CIA + 4.
<b>bnl</b>		Branch if not less than; use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+0,target</b>	
<b>bnla</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+0,target</b>	
<b>bnll</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+0,target</b>	(LR) ← CIA + 4.
<b>bnlla</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+0,target</b>	(LR) ← CIA + 4.
<b>bns</b>		Branch if not summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+3,target</b>	
<b>bnsa</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+3,target</b>	
<b>bnsl</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bnsla</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+3,target</b>	(LR) ← CIA + 4.

Table 9-8. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bnu</b>		Branch if not unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+3,target</b>	
<b>bnua</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+3,target</b>	
<b>bnul</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bnula</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bso</b>		Branch if summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+3,target</b>	
<b>bsoa</b>	[cr_field,] target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+3,target</b>	
<b>bsol</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bsola</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bt</b>		Branch if CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 12,cr_bit,target</b>	
<b>bta</b>	cr_bit, target	<i>Extended mnemonic for</i> <b>bca 12,cr_bit,target</b>	
<b>btl</b>		<i>Extended mnemonic for</i> <b>bcl 12,cr_bit,target</b>	(LR) ← CIA + 4.
<b>btla</b>		<i>Extended mnemonic for</i> <b>bcla 12,cr_bit,target</b>	(LR) ← CIA + 4.
<b>bun</b>		Branch if unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+3,target</b>	
<b>buna</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+3,target</b>	
<b>bunl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+3,target</b>	(LR) ← CIA + 4.
<b>bunla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+3,target</b>	(LR) ← CIA + 4.

## bcctr

Branch Conditional to Count Register  
PPC440x5 CPU Core User's Manual



Preliminary

**bcctr** BO, BI LK=0  
**bcctrl** BO, BI LK=1

19	BO	BI		528	LK
0	6	11	16	21	31

```
if (BO0 = 1 ∨ (CRBI = BO1)) then
    NIA ← CTR0:29 || 20
else
    NIA ← CIA + 4
if LK = 1 then
    (LR) ← CIA + 4
PC ← NIA
```

If BO<sub>0</sub> contains 0, then the CR bit specified by the BI field is compared to BO<sub>1</sub> as part of the branch condition. If BO<sub>0</sub> contains 1, then the CR is not tested as part of the branch condition, and the BI field is ignored.

The next instruction address (NIA) is either the effective address of the branch target, or the address of the instruction after the branch, depending on whether the branch is taken or not. The branch target address is formed by concatenating two 0-bits to the right of the 30 most significant bits of the CTR.

BO<sub>4</sub> affects branch prediction, a performance-improvement feature. See *Branch Prediction* on page 64 for a complete discussion.

Instruction execution resumes with the instruction at the NIA.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

### Registers Altered

- LR if LK contains 1

### Invalid Instruction Forms

- Reserved fields
- If BO<sub>2</sub> contains 0, the instruction form is invalid, and the result of the instruction (in particular, the branch target address and whether or not the branch is taken) is undefined. The architecture does not permit the combination of decrementing the CTR as part of the branch condition, together with using the CTR as the branch target address.

Table 9-9. Extended Mnemonics for *bcctr*, *bcctrl*

Mnemonic	Operands	Function	Other Registers Altered
<b>bcctr</b>		Branch unconditionally to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 20,0</b>	
<b>bcctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 20,0</b>	(LR) ← CIA + 4.

Table 9-9. Extended Mnemonics for bcctr, bcctrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>beqctr</b>	[cr_field]	Branch, if equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+2</b>	
<b>beqctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+2</b>	(LR) ← CIA + 4.
<b>bfctr</b>	cr_bit	Branch, if CR <sub>cr_bit</sub> = 0, to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 4,cr_bit</b>	
<b>bfctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,cr_bit</b>	(LR) ← CIA + 4.
<b>bgectr</b>	[cr_field]	Branch, if greater than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+0</b>	
<b>bgectrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bgctr</b>	[cr_field]	Branch, if greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+1</b>	
<b>bgctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+1</b>	(LR) ← CIA + 4.
<b>blectr</b>	[cr_field]	Branch, if less than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+1</b>	
<b>blectrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+1</b>	(LR) ← CIA + 4.
<b>bltctr</b>	[cr_field]	Branch, if less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+0</b>	
<b>bltctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bnctr</b>	[cr_field]	Branch, if not equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+2</b>	
<b>bnctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+2</b>	(LR) ← CIA + 4.
<b>bngctr</b>	[cr_field]	Branch, if not greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+1</b>	
<b>bngctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+1</b>	(LR) ← CIA + 4.

# bcctr

Branch Conditional to Count Register  
PPC440x5 CPU Core User's Manual




Preliminary

Table 9-9. Extended Mnemonics for bcctr, bcctrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bnlctr</b>	[cr_field]	Branch, if not less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+0</b>	
<b>bnlctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bnsctr</b>	[cr_field]	Branch, if not summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+3</b>	
<b>bnsctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>bnuctr</b>	[cr_field]	Branch, if not unordered, to address in CTR; use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+3</b>	
<b>bnuctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>bsocctr</b>	[cr_field]	Branch, if summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+3</b>	
<b>bsocctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>btctr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 1 to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 12,cr_bit</b>	
<b>btctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,cr_bit</b>	(LR) ← CIA + 4.
<b>bunctr</b>	[cr_field]	Branch if unordered to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+3</b>	
<b>bunctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+3</b>	(LR) ← CIA + 4.

**Preliminary**

<b>bclr</b>	BO, BI	LK=0
<b>bclrl</b>	BO, BI	LK=1



```

if BO2 = 0 then
  CTR ← CTR - 1
if (BO2 = 1 ∨ ((CTR = 0) = BO3)) ∧ (BO0 = 1 ∨ (CRBI = BO1)) then
  NIA ← LR0:29 || 20
else
  NIA ← CIA + 4
if LK = 1 then
  (LR) ← CIA + 4
PC ← NIA
  
```

If BO<sub>2</sub> contains 0, the CTR decrements, and the decremented value is tested for 0 as part of the branch condition. In this case, BO<sub>3</sub> indicates whether the test for 0 must be true or false in order for the branch to be taken. If BO<sub>2</sub> contains 1, then the CTR is neither decremented nor tested as part of the branch condition.

If BO<sub>0</sub> contains 0, then the CR bit specified by the BI field is compared to BO<sub>1</sub> as part of the branch condition. If BO<sub>0</sub> contains 1, then the CR is not tested as part of the branch condition, and the BI field is ignored.

The next instruction address (NIA) is either the effective address of the branch target, or the address of the instruction after the branch, depending on whether the branch is taken or not. The branch target address is formed by concatenating two 0-bits to the right of the 30 most significant bits of the LR.

BO<sub>4</sub> affects branch prediction, a performance-improvement feature. See *Branch Prediction* on page 64 for a complete discussion.

Instruction execution resumes with the instruction at the NIA.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

**Registers Altered**

- CTR if BO<sub>2</sub> contains 0
- LR if LK contains 1

**Invalid Instruction Forms**

- Reserved fields

Table 9-10. Extended Mnemonics for bclr, bclrl

Mnemonic	Operands	Function	Other Registers Altered
<b>blr</b>		Branch unconditionally to address in LR. <i>Extended mnemonic for</i> <b>bclr 20,0</b>	
<b>bclrl</b>		<i>Extended mnemonic for</i> <b>bclrl 20,0</b>	(LR) ← CIA + 4.
<b>bdnzlr</b>		Decrement CTR. Branch if CTR ≠ 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 16,0</b>	
<b>bdnzlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 16,0</b>	(LR) ← CIA + 4.
<b>bdnzflr</b>	cr_bit	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 0,cr_bit</b>	
<b>bdnzflrl</b>		<i>Extended mnemonic for</i> <b>bclrl 0,cr_bit</b>	(LR) ← CIA + 4.
<b>bdnztlr</b>	cr_bit	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 1 to address in LR. <i>Extended mnemonic for</i> <b>bclr 8,cr_bit</b>	
<b>bdnztlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 8,cr_bit</b>	(LR) ← CIA + 4.
<b>bdzlr</b>		Decrement CTR. Branch if CTR = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 18,0</b>	
<b>bdzlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 18,0</b>	(LR) ← CIA + 4.
<b>bdzflr</b>	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 2,cr_bit</b>	
<b>bdzflrl</b>		<i>Extended mnemonic for</i> <b>bclrl 2,cr_bit</b>	(LR) ← CIA + 4.
<b>bdztlr</b>	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 1 to address in LR. <i>Extended mnemonic for</i> <b>bclr 10,cr_bit</b>	
<b>bdztlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 10,cr_bit</b>	(LR) ← CIA + 4.



Table 9-10. Extended Mnemonics for bclr, bclrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bqlr</b>	[cr_field]	Branch if equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+2</b>	
<b>bqlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+2</b>	(LR) ← CIA + 4.
<b>bflr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 4,cr_bit</b>	
<b>bflrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,cr_bit</b>	(LR) ← CIA + 4.
<b>bgelr</b>	[cr_field]	Branch, if greater than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+0</b>	
<b>bgelrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bgtlr</b>	[cr_field]	Branch, if greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+1</b>	
<b>bgtlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+1</b>	(LR) ← CIA + 4.
<b>blelr</b>	[cr_field]	Branch, if less than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+1</b>	
<b>blelrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+1</b>	(LR) ← CIA + 4.
<b>bltlr</b>	[cr_field]	Branch, if less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+0</b>	
<b>bltlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bnelr</b>	[cr_field]	Branch, if not equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+2</b>	
<b>bnelrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+2</b>	(LR) ← CIA + 4.
<b>bnglr</b>	[cr_field]	Branch, if not greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+1</b>	
<b>bnglrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+1</b>	(LR) ← CIA + 4.


**bclr**Branch Conditional to Link Register  
PPC440x5 CPU Core User's Manual

Preliminary

Table 9-10. Extended Mnemonics for *bclr*, *bclrl* (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>bnlrl</b>	[cr_field]	Branch, if not less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+0</b>	
<b>bnlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+0</b>	(LR) ← CIA + 4.
<b>bnsrl</b>	[cr_field]	Branch if not summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+3</b>	
<b>bnsrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>bnulrl</b>	[cr_field]	Branch if not unordered to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+3</b>	
<b>bnulrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>bsolrl</b>	[cr_field]	Branch if summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+3</b>	
<b>bsolrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+3</b>	(LR) ← CIA + 4.
<b>btlrl</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 1 to address in LR. <i>Extended mnemonic for</i> <b>bclr 12,cr_bit</b>	
<b>btlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,cr_bit</b>	(LR) ← CIA + 4.
<b>bunlrl</b>	[cr_field]	Branch if unordered to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+3</b>	
<b>bunlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+3</b>	(LR) ← CIA + 4.

**cmp** BF, 0, RA, RB



```

c0:3 ← 40
if (RA) < (RB) then c0 ← 1
if (RA) > (RB) then c1 ← 1
if (RA) = (RB) then c2 ← 1
c3 ← XER[SO]
n ← BF
CR[CRn] ← c0:3
    
```

The contents of register RA are compared with the contents of register RB using a 32-bit signed compare.

The CR field specified by the BF field is updated to reflect the results of the compare and the value of XER[SO] is placed into the same CR field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- CR[CRn] where n is specified by the BF field

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

PowerPC Book-E architecture defines this instruction as **cmp BF,L,RA,RB**, where L selects operand size for 64-bit implementations. For all 32-bit implementations, L = 0 is required (L = 1 is an invalid form); hence for the PPC440x5 core, use of the extended mnemonic **cmpw BF,RA,RB** is recommended.

*Table 9-11. Extended Mnemonics for cmp*

Mnemonic	Operands	Function	Other Registers Altered
<b>cmpw</b>	[BF,] RA, RB	Compare Word; use CR0 if BF is omitted. <i>Extended mnemonic for</i> <b>cmp BF,0,RA,RB</b>	

# cmpi


Compare Immediate

PPC440x5 CPU Core User's Manual



Preliminary

**cmpi** BF, 0, RA, IM



$c_{0:3} \leftarrow 4_0$   
if (RA) < EXTS(IM) then  $c_0 \leftarrow 1$   
if (RA) > EXTS(IM) then  $c_1 \leftarrow 1$   
if (RA) = EXTS(IM) then  $c_2 \leftarrow 1$   
 $c_3 \leftarrow \text{XER}[\text{SO}]$   
 $n \leftarrow \text{BF}$   
 $\text{CR}[\text{CR}_n] \leftarrow c_{0:3}$

The IM field is sign-extended to 32 bits. The contents of register RA are compared with the extended IM field, using a 32-bit signed compare.

The CR field specified by the BF field is updated to reflect the results of the compare and the value of XER[SO] is placed into the same CR field.

## Registers Altered

- CR[CR<sub>n</sub>] where *n* is specified by the BF field

## Invalid Instruction Forms

- Reserved fields


## Programming Note

PowerPC Book-E Architecture defines this instruction as **cmpi BF,L,RA,IM**, where L selects operand size for 64-bit implementations. For all 32-bit implementations, L = 0 is required (L = 1 is an invalid form); hence for the PPC440x5 core, use of the extended mnemonic **cmpwi BF,RA,IM** is recommended.

Table 9-12. Extended Mnemonics for cmpi

Mnemonic	Operands	Function	Other Registers Altered
<b>cmpwi</b>	[BF,] RA, IM	Compare Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> <b>cmpi BF,0,RA,IM</b>	

**cmpl** BF, 0, RA, RB



```

c0:3 ← 40
if (RA) <u (RB) then c0 ← 1
if (RA) >u (RB) then c1 ← 1
if (RA) = (RB) then c2 ← 1
c3 ← XER[SO]
n ← BF
CR[CRn] ← c0:3
    
```

The contents of register RA are compared with the contents of register RB, using a 32-bit unsigned compare.

The CR field specified by the BF field is updated to reflect the results of the compare and the value of XER[SO] is placed into the same CR field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- CR[CRn] where n is specified by the BF field

**Invalid Instruction Forms**

- Reserved fields

**Programming Notes**

PowerPC Book-E Architecture defines this instruction as **cmpl BF,L,RA,RB**, where L selects operand size for 64-bit implementations. For all 32-bit implementations, L = 0 is required (L = 1 is an invalid form); hence for PPC440x5 core, use of the extended mnemonic **cmplw BF,RA,RB** is recommended.

*Table 9-13. Extended Mnemonics for cmpl*

Mnemonic	Operands	Function	Other Registers Altered
<b>cmplw</b>	[BF,] RA, RB	Compare Logical Word. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> <b>cmpl BF,0,RA,RB</b>	


# cmpli

Compare Logical Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

**cmpli** BF, 0, RA, IM



```

c0:3 ← 40
if (RA) < (160 || IM) then c0 ← 1
if (RA) > (160 || IM) then c1 ← 1
if (RA) = (160 || IM) then c2 ← 1
c3 ← XER[SO]
n ← BF
CR[CRn] ← c0:3

```

The IM field is extended to 32 bits by concatenating 16 0-bits to its left. The contents of register RA are compared with IM using a 32-bit unsigned compare.

The CR field specified by the BF field is updated to reflect the results of the compare and the value of XER[SO] is placed into the same CR field.

### Registers Altered

- CR[CRn] where n is specified by the BF field

### Invalid Instruction Forms

- Reserved fields

### Programming Note


PowerPC Book-E Architecture defines this instruction as **cmpli BF,L,RA,IM**, where L selects operand size for 64-bit implementations. For all 32-bit implementations, L = 0 is required (L = 1 is an invalid form); hence for the PPC440x5 core, use of the extended mnemonic **cmplwi BF,RA,IM** is recommended.

Table 9-14. Extended Mnemonics for *cmpli*

Mnemonic	Operands	Function	Other Registers Changed
<b>cmplwi</b>	[BF,] RA, IM	Compare Logical Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> <b>cmpli BF,0,RA,IM</b>	



**cntlzw** RA, RS Rc=0  
**cntlzw.** RA, RS Rc=1



```
n ← 0
do while n < 32
  if (RS)n = 1 then leave
  n ← n + 1
(RA) ← n
```

The consecutive leading 0 bits in register RS are counted; the count is placed into register RA.

The count ranges from 0 through 32, inclusive.

**Registers Altered**

- RA
- CR[CR0] if Rc contains 1

**Invalid Instruction Forms**

- Reserved fields

## crand


Condition Register AND

PPC440x5 CPU Core User's Manual



Preliminary

**crand**      BT, BA, BB



$$CR_{BT} \leftarrow CR_{BA} \wedge CR_{BB}$$

The CR bit specified by the BA field is ANDed with the CR bit specified by the BB field; the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- $CR_{BT}$


### Invalid Instruction Forms

- Reserved fields





**crandc**      BT, BA, BB



$$CR_{BT} \leftarrow CR_{BA} \wedge \neg CR_{BB}$$

The CR bit specified by the BA field is ANDed with the ones complement of the CR bit specified by the BB field; the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- $CR_{BT}$

**Invalid Instruction Forms**

- Reserved fields


# creqv

Condition Register Equivalent  
PPC440x5 CPU Core User's Manual



Preliminary

**creqv** BT, BA, BB



$$CR_{BT} \leftarrow \neg(CR_{BA} \oplus CR_{BB})$$

The CR bit specified by the BA field is XORed with the CR bit specified by the BB field; the ones complement of the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- $CR_{BT}$

### Invalid Instruction Forms


- Reserved fields

Table 9-15. Extended Mnemonics for creqv

Mnemonic	Operands	Function	Other Registers Altered
crset	bx	CR set. <i>Extended mnemonic for creqv bx,bx,bx</i>	



**crnand**      BT, BA, BB



$$CR_{BT} \leftarrow \neg(CR_{BA} \wedge CR_{BB})$$

The CR bit specified by the BA field is ANDed with the CR bit specified by the BB field; the ones complement of the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- $CR_{BT}$

**Invalid Instruction Forms**

- Reserved fields

## crnor


Condition Register NOR

PPC440x5 CPU Core User's Manual



Preliminary

**crnor**      BT, BA, BB



$$CR_{BT} \leftarrow \neg(CR_{BA} \vee CR_{BB})$$

The CR bit specified by the BA field is ORed with the CR bit specified by the BB field; the ones complement of the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- $CR_{BT}$


### Invalid Instruction Forms

- Reserved fields

Table 9-16. Extended Mnemonics for crnor

Mnemonic	Operands	Function	Other Registers Altered
<b>crnot</b>	bx, by	CR not. <i>Extended mnemonic for crnor bx,by,by</i>	

**cror**            BT, BA, BB



$$CR_{BT} \leftarrow CR_{BA} \vee CR_{BB}$$

The CR bit specified by the BA field is ORed with the CR bit specified by the BB field; the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- CR<sub>BT</sub>

**Invalid Instruction Forms**

- Reserved fields

*Table 9-17. Extended Mnemonics for cror*

Mnemonic	Operands	Function	Other Registers Altered
<b>crmove</b>	bx, by	CR move. <i>Extended mnemonic for cror bx,by,by</i>	


## crorc

Condition Register OR with Complement  
PPC440x5 CPU Core User's Manual



Preliminary

**crorc**            BT, BA, BB



$$CR_{BT} \leftarrow CR_{BA} \vee \neg CR_{BB}$$

The condition register (CR) bit specified by the BA field is ORed with the ones complement of the CR bit specified by the BB field; the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered


- $CR_{BT}$

### Invalid Instruction Forms

- Reserved fields

Preliminary

**crxor**      BT, BA, BB



$$CR_{BT} \leftarrow CR_{BA} \oplus CR_{BB}$$

The CR bit specified by the BA field is XORed with the CR bit specified by the BB field; the result is placed into the CR bit specified by the BT field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- CR<sub>BT</sub>

**Invalid Instruction Forms**

- Reserved fields

*Table 9-18. Extended Mnemonics for crxor*

Mnemonic	Operands	Function	Other Registers Altered
<b>crclr</b>	bx	Condition register clear. <i>Extended mnemonic for crxor bx,bx,bx</i>	


# dcba

Data Cache Block Allocate  
PPC440x5 CPU Core User's Manual



Preliminary


**dcba** RA, RB



**dcba** is treated as a no-op by the PPC440x5 core.



**dcbf** RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$DCBF(EA)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the data block corresponding to the EA is in the data cache and marked as modified (stored into), the data block is copied back to main storage and then marked invalid in the data cache. If the data block is not marked as modified, it is simply marked invalid in the data cache. The operation is performed whether or not the memory page referenced by the EA is marked as cacheable.

If the data block at the EA is not in the data cache, no operation is performed.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Exceptions**

This instruction is considered a “load” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “store” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.

This instruction may cause a Cache Locking type of Data Storage exception. See *Data Storage Interrupt* on page 175 for more information.


## dcbi

Data Cache Block Invalidate  
PPC440x5 CPU Core User's Manual



Preliminary

**dcbi** RA, RB



$EA \leftarrow (RA|0) + (RB)$   
DCBI(EA)

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the data block at the EA is in the data cache, the data block is marked invalid, regardless of whether or not the memory page referenced by the EA is marked as cacheable. If modified data existed in the data block prior to the operation of this instruction, that data is lost.

If the data block at the EA is not in the data cache, no operation is performed.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Notes


Execution of this instruction is privileged.

### Exceptions

This instruction is considered a “store” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “store” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.

**dcbst**      RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$DCBST(EA)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0, and is the contents of register RA otherwise.

If the data block at the EA is in the data cache and marked as modified, the data block is copied back to main storage and marked as unmodified in the data cache.

If the data block at the EA is in the data cache, and is not marked as modified, or if the data block at the EA is not in the data cache, no operation is performed.

The operation specified by this instruction is performed whether or not the memory page referenced by the EA is marked as cacheable.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Exceptions**

This instruction is considered a “load” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “store” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.


## dcbt

Data Cache Block Touch  
PPC440x5 CPU Core User's Manual



Preliminary

**dcbt** RA, RB



$EA \leftarrow (RA|0) + (RB)$   
DCBT(EA)

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

If the data block at the EA is not in the data cache and the memory page referenced by the EA is marked as cacheable, the block is read from main storage into the data cache.

If the data block at the EA is in the data cache, or if the memory page referenced by the EA is marked as caching inhibited, no operation is performed.

This instruction is not allowed to cause Data Storage interrupts nor Data TLB Error interrupts. If execution of the instruction causes either of these types of exception, then no operation is performed, and no interrupt occurs.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields


### Programming Notes

The **dcbt** instruction allows a program to begin a cache block fetch from main storage before the program needs the data. The program can later load data from the cache into registers without incurring the latency of a cache miss.

### Exceptions

This instruction is considered a “load” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “load” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.

**dcbtst** RA, RB


$$EA \leftarrow (RA|0) + (RB)$$

$$DCBTST(EA)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the data block at the EA is not in the data cache and the memory page referenced by the EA address is marked as cacheable, the data block is loaded into the data cache.

If the data block at the EA is in the data cache, or if the memory page referenced by the EA is marked as caching inhibited, no operation is performed.

This instruction is not allowed to cause Data Storage interrupts nor Data TLB Error interrupts. If execution of the instruction causes either of these types of exception, then no operation is performed, and no interrupt occurs.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

The **dcbtst** instruction allows a program to begin a cache block fetch from main storage before the program needs the data. The program can later store data from GPRs into the cache block, without incurring the latency of a cache miss.

Architecturally, **dcbtst** is intended to bring a cache block into the data cache in a manner which will permit future instructions to store to that block efficiently. For example, in an implementation which supports the "MESI" cache coherency protocol, the block would be brought into the cache in "Exclusive" mode, allowing the block to be stored to without having to broadcast any coherency operations on the system bus. However, since the PPC440x5 core does not support hardware-enforcement of multiprocessor coherency, there is no distinction between a block being brought in for a read or a write, and hence the implementation of the **dcbtst** instruction is identical to the implementation of the **dcbt** instruction.

### Exceptions

This instruction is considered a "load" with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a "load" with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.


## dcbz

Data Cache Block Set to Zero  
PPC440x5 CPU Core User's Manual



Preliminary

**dcbz** RA, RB



$EA \leftarrow (RA|0) + (RB)$   
DCBZ(EA)

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the data block at the EA is in the data cache and the memory page referenced by the EA is marked as cacheable and non-write-through, the data in the cache block is set to 0 and marked as *dirty* (modified).

If the data block at the EA is not in the data cache and the memory page referenced by the EA is marked as cacheable and non-write-through, a cache block is established and set to 0 and marked as dirty. Note that nothing is read from main storage, as described in the programming note.

If the memory page referenced by the EA is marked as either write-through or as caching inhibited, an Alignment exception occurs.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

Because **dcbz** can establish an address in the data cache without copying the contents of that address from main storage, the address established may be invalid with respect to the storage subsystem. A subsequent operation may cause the address to be copied back to main storage, for example, to make room for a new cache block; a Data Machine Check exception could occur under these circumstances.

If **dcbz** is attempted to an EA in a memory page which is marked as caching inhibited or as write-through, the software alignment exception handler should emulate the instruction by storing zeros to the block referenced by the EA. The store instructions in the emulation software will cause main storage to be updated (and possibly the cache, if the EA is in a page marked as write-through).

### Exceptions


An alignment exception occurs if the EA is marked as caching inhibited or as write-through.

This instruction is considered a “store” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “store” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.



**dccci** RA, RB



**DCCCI**

This instruction flash invalidates the entire data cache array. The RA and RB operands are not used; previous implementations used these operands to calculate an effective address (EA) which specified the particular block or blocks to be invalidated. The instruction form (including the specification of RA and RB operands) is maintained for software and tool compatibility.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Programming Notes**

Execution of this instruction is privileged.

This instruction is intended for use in the power-on reset routine to invalidate the entire data cache array before caching is enabled.


**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

# dcread



**dcread** RT, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $INDEX \leftarrow EA_{17:26}$   
 $WORD \leftarrow EA_{27:29}$   
 $(RT) \leftarrow (\text{data cache data})[INDEX,WORD]$   
 $DCDBTRH \leftarrow (\text{data cache tag high})[INDEX]$   
 $DCDBTRL \leftarrow (\text{data cache tag low})[INDEX]$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

EA<sub>17:26</sub> selects a line of tag and data from the data cache. EA<sub>27:29</sub> selects a word from the 8-word data portion of the selected cache line, and this word is read into register RT. EA<sub>30:31</sub> must be 0b00; if not, the value placed in register RT is undefined.

The tag portion of the selected cache line is read into the DCDBTRH and DCDBTRL registers, as follows:

Register[bit(s)]	Tag Field	Name	
DCDBTRH[0:23]	TRA	Tag Real Address	Bits 0:23 of the lower 32 bits of the 36-bit real address associated with this cache line
DCDBTRH[24]	V	Valid	The valid indicator for the cache line (1 indicates valid)
DCDBTRH[25:27]		reserved	Reserved fields are read as 0s
DCDBTRH[28:31]	TERA	Tag Extended Real Address	Upper 4 bits of the 36-bit real address associated with this cache line
DCDBTRL[0:23]		reserved	Reserved fields are read as 0s
DCDBTRL[24:27]	D	Dirty Indicators	The "dirty" (modified) indicators for each of the four doublewords in the cache line
DCDBTRL[28]	U0	U0 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line
DCDBTRL[29]	U1	U1 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line
DCDBTRL[30]	U2	U2 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line
DCDBTRL[31]	U3	U3 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line

This instruction can be used by a debug tool to determine the contents of the data cache, without knowing the specific addresses of the lines which are currently contained within the cache.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.



### Registers Altered

- RT
- DCDBTRH
- DCDBTRL

### Invalid Instruction Forms

- Reserved fields

### Programming Note

Execution of this instruction is privileged.

The PPC440x5 core does not support the use of the **dcread** instruction when the data cache controller is still in the process of performing cache operations associated with previously executed instructions (such as line fills and line flushes). Also, the PPC440x5 core does not automatically synchronize context between a **dcread** instruction and the subsequent **mfspr** instructions that read the results of the **dcread** instruction into GPRs. In order to guarantee that the **dcread** instruction operates correctly, and that the **mfspr** instructions obtain the results of the **dcread** instruction, a sequence such as the following must be used:

```

msync                # ensure that all previous cache operations have completed
dcread    regT,regA,regB # read cache information; the contents of GPR A and GPR B are
                        # added and the result used to specify a cache line index to be read;
                        # the data word is moved into GPR T and the tag information is read
                        # into DCDBTRH and DCDBTRL

isync                # ensure dcread completes before attempting to read results
mfdcdbtrh    regD    # move high portion of tag into GPR D
mfdcdbtrl    regE    # move low portion of tag into GPR E

```

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## divw


Divide Word

PPC440x5 CPU Core User's Manual



Preliminary

<b>divw</b>	RT, RA, RB	OE=0, Rc=0
<b>divw.</b>	RT, RA, RB	OE=0, Rc=1
<b>divwo</b>	RT, RA, RB	OE=1, Rc=0
<b>divwo.</b>	RT, RA, RB	OE=1, Rc=1



$$(RT) \leftarrow (RA) \div (RB)$$

The contents of register RA are divided by the contents of register RB. The quotient is placed into register RT.

Both the dividend and the divisor are interpreted as signed integers. The quotient is the unique signed integer that satisfies:

$$\text{dividend} = (\text{quotient} \times \text{divisor}) + \text{remainder}$$

where the remainder has the same sign as the dividend and its magnitude is less than that of the divisor.

If an attempt is made to perform  $(0x8000\ 0000 \div -1)$  or  $(n \div 0)$ , the contents of register RT are undefined; if the R<sub>c</sub> field also contains 1, the contents of CR[CR0]<sub>0:2</sub> are undefined. Either invalid division operation sets XER[OV, SO] (and CR[CR0]<sub>3</sub> if R<sub>c</sub> contains 1) to 1 if the OE field contains 1.

### Registers Altered

- RT
- CR[CR0] if R<sub>c</sub> contains 1
- XER[OV, SO] if OE contains 1


### Programming Note

The 32-bit remainder can be calculated using the following sequence of instructions:

divw	RT,RA,RB	# RT = quotient
mullw	RT,RT,RB	# RT = quotient × divisor
subf	RT,RT,RA	# RT = remainder

The sequence does not calculate correct results for the invalid divide operations.

<b>divwu</b>	RT, RA, RB	OE=0, Rc=0
<b>divwu.</b>	RT, RA, RB	OE=0, Rc=1
<b>divwuo</b>	RT, RA, RB	OE=1, Rc=0
<b>divwuo.</b>	RT, RA, RB	OE=1, Rc=1



$$(RT) \leftarrow (RA) \div (RB)$$

The contents of register RA are divided by the contents of register RB. The quotient is placed into register RT.

The dividend and the divisor are interpreted as unsigned integers. The quotient is the unique unsigned integer that satisfies:

$$\text{dividend} = (\text{quotient} \times \text{divisor}) + \text{remainder}$$

If an attempt is made to perform  $(n \div 0)$ , the contents of register RT are undefined; if the Rc also contains 1, the contents of CR[CR0]<sub>0,2</sub> are also undefined. The invalid division operation also sets XER[OV, SO] (and CR[CR0]<sub>3</sub> if Rc contains 1) to 1 if the OE field contains 1.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[OV, SO] if OE contains 1

**Programming Note**

The 32-bit remainder can be calculated using the following sequence of instructions

```
divwu    RT,RA,RB      # RT = quotient
mullw   RT,RT,RB      # RT = quotient × divisor
subf    RT,RT,RA      # RT = remainder
```

This sequence does not calculate the correct result if the divisor is 0.


## dImzb

Determine Leftmost Zero Byte  
PPC440x5 CPU Core User's Manual



Preliminary

<b>dImzb</b>	RA, RS, RB	Rc=0
<b>dImzb.</b>	RA, RS, RB	Rc=1



```
d ← (RS) || (RB)
i, x, y ← 0
do while (x < 8) ∧ (y = 0)
  x ← x + 1
  if di:i+7 = 0 then
    y ← 1
  else
    i ← i + 8
(RA) ← x
XER[TBC] ← x
if Rc = 1 then
  CR[CR0]3 ← XER[SO]
  if y = 1 then
    if x < 5 then
      CR[CR0]0:2 ← 0b010
    else
      CR[CR0]0:2 ← 0b100
  else
    CR[CR0]0:2 ← 0b001
```

The contents of registers RS and RB are concatenated to form an 8-byte operand. The operand is searched for the leftmost byte in which each bit is 0 (a 0-byte).

Bytes in the operand are numbered from left to right starting with 1. If a 0-byte is found, its byte number is placed into XER[TBC] and register RA. Otherwise, the number 8 is placed into XER[TBC] and register RA.

If the Rc field contains 1, XER[SO] is copied to CR[CR0]<sub>3</sub> and CR[CR0]<sub>0:2</sub> are updated as follows:

- If no 0-byte is found, CR[CR0]<sub>0:2</sub> is set to 0b001.
- If the leftmost 0-byte is in the first 4 bytes (in the RS register), CR[CR0]<sub>0:2</sub> is set to 0b010.
- If the leftmost 0-byte is in the last 4 bytes (in the RB register), CR[CR0]<sub>0:2</sub> is set to 0b100.

### Registers Altered


- XER[TBC]
- RA
- CR[CR0] if Rc contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**eqv** RA, RS, RB R<sub>c</sub>=0  
**eqv.** RA, RS, RB R<sub>c</sub>=1



$$(RA) \leftarrow \neg((RS) \oplus (RB))$$

The contents of register RS are XORed with the contents of register RB; the ones complement of the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if R<sub>c</sub> contains 1

## extsb


Extend Sign Byte

PPC440x5 CPU Core User's Manual



Preliminary

<b>extsb</b>	RA, RS	Rc=0
<b>extsb.</b>	RA, RS	Rc=1



$(RA) \leftarrow \text{EXTS}(RS)_{24:31}$

The least significant byte of register RS is sign-extended to 32 bits by replicating bit 24 of the register into bits 0 through 23 of the result. The result is placed into register RA.

### Registers Altered


- RA
- CR[CR0] if Rc contains 1

### Invalid Instruction Forms

- Reserved fields



**extsh** RA, RS Rc=0  
**extsh.** RA, RS Rc=1



$$(RA) \leftarrow \text{EXTS}(RS)_{16:31}$$

The least significant halfword of register RS is sign-extended to 32 bits by replicating bit 16 of the register into bits 0 through 15 of the result. The result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if Rc contains 1

**Invalid Instruction Forms**

- Reserved fields


## icbi

Instruction Cache Block Invalidate  
PPC440x5 CPU Core User's Manual



Preliminary

**icbi** RA, RB



$EA \leftarrow (RA|0) + (RB)$   
ICBI(EA)

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the instruction block at the EA is in the instruction cache, the cache block is marked invalid.

If the instruction block at the EA is not in the instruction cache, no additional operation is performed.

The operation specified by this instruction is performed whether or not the EA is marked as cacheable.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Note

Instruction cache management instructions use MSR[DS], not MSR[IS], as part of the virtual address. Also, the instruction cache on the PPC440x5 is “virtually-tagged”, which means that the EA is converted to a virtual address (VA), and the VA is compared against the cache tag field. See *Instruction Cache Synonyms* on page 104 for more information on the ramifications of virtual tagging on software.

### Exceptions


Instruction Storage interrupts and Instruction TLB Error interrupts are associated with exceptions which occur during instruction *fetching*, not during instruction *execution*. *Execution* of instruction cache management instructions may cause Data Storage or Data TLB Error exceptions.

This instruction is considered a “load” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “load” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.

This instruction may cause a Cache Locking type of Data Storage exception. See *Data Storage Interrupt* on page 175 for more information.



**icbt** RA, RB


$$EA \leftarrow (RA|0) + (RB)$$

$$ICBT(EA)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

If the instruction block at the EA is not in the instruction cache and the memory page referenced by the EA is marked as cacheable, the instruction block is fetched into the instruction cache.

If the instruction block at the EA is in the instruction cache, or if the memory page referenced by the EA is marked as caching inhibited, no operation is performed.

If the memory page referenced by the EA is marked as “no-execute” for the current operating mode (user mode or supervisor mode, as specified by MSR[PR]), no operation is performed.

This instruction is not allowed to cause Data Storage interrupts nor Data TLB Error interrupts. If execution of the instruction causes either of these types of exception, then no operation is performed, and no interrupt occurs.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

This instruction allows a program to begin a cache block fetch from main storage before the program needs the instruction. The program can later branch to the instruction address and fetch the instruction from the cache without incurring the latency of a cache miss.

Instruction cache management instructions use MSR[DS], not MSR[IS], as part of the virtual address. Also, the instruction cache on the PPC440x5 is “virtually-tagged”, which means that the EA is converted to a virtual address (VA), and the VA is compared against the cache tag field. See *Instruction Cache Synonyms* on page 104 for more information on the ramifications of virtual tagging on software.

### Exceptions

Instruction Storage interrupts and Instruction TLB Error interrupts are associated with exceptions which occur during instruction *fetching*, not during instruction *execution*. *Execution* of instruction cache management instructions may cause Data Storage or Data TLB Error exceptions, but are not allowed to cause the associated interrupt. Instead, if such an exception occurs, then no operation is performed.

## icbt

Instruction Cache Block Touch  
PPC440x5 CPU Core User's Manual




Preliminary

---

This instruction is considered a “load” with respect to Data Storage exceptions. See *Data Storage Interrupt* on page 175 for more information.

This instruction is considered a “load” with respect to data address compare (DAC) Debug exceptions. See *Debug Interrupt* on page 188 for more information.

**iccci**      RA, RB

**ICCCI**

This instruction flash invalidates the entire instruction cache array. The RA and RB operands are not used; previous implementations used these operands to calculate an effective address (EA) which specified the particular block or blocks to be invalidated. The instruction form (including the specification of RA and RB operands) is maintained for software and tool compatibility.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Programming Notes**

Execution of this instruction is privileged.

This instruction is intended for use in the power-on reset routine to invalidate the entire instruction cache array before caching is enabled.

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.


## icread

Instruction Cache Read  
PPC440x5 CPU Core User's Manual



Preliminary

**icread** RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $INDEX \leftarrow EA_{17:26}$   
 $WORD \leftarrow EA_{27:29}$   
 $ICDBDR \leftarrow (\text{instruction cache data})[INDEX,WORD]$   
 $ICDBTRH \leftarrow (\text{instruction cache tag high})[INDEX]$   
 $ICDBTRL \leftarrow (\text{instruction cache tag low})[INDEX]$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

$EA_{17:26}$  selects a line of tag and data (instructions) from the instruction cache.  $EA_{27:29}$  selects a 32-bit instruction from the 8-instruction data portion of the selected cache line, and this instruction is read into the ICDBDR.  $EA_{30:31}$  are ignored, as are  $EA_{0:16}$ .

The tag portion of the selected cache line is read into the ICDBTRH and ICDBTRL registers, as follows:

Register[bit(s)]	Tag Field	Name	
ICDBTRH[0:23]	TEA	Tag Effective Address	Bits 0:23 of the 32-bit effective address associated with this cache line
ICDBTRH[24]	V	Valid	The valid indicator for the cache line (1 indicates valid)
ICDBTRH[25:31]		reserved	Reserved fields are read as 0s
ICDBTRL[0:21]		reserved	Reserved fields are read as 0s
ICDBTRL[22]	TS	Translation Space	The address space portion of the virtual address associated with this cache line.
ICDBTRL[23]	TD	Translation ID (TID) Disable	TID Disable field for the memory page associated with this cache line
ICDBTRL[24:31]	TID	Translation ID	TID field portion of the virtual address associated with this cache line

The instruction cache on PPC440x5 is “virtually-tagged”, which means that the tag field contains the virtual address, which consists of the TEA, TS, and TID fields. See *Memory Management* on page 129 for more information on the function of the TS, TD, and TID fields.

This instruction can be used by a debug tool to determine the contents of the instruction cache, without knowing the specific addresses of the lines which are currently contained within the cache.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Preliminary

**Registers Altered**

- ICDBDR
- ICDBTRH
- ICDBTRL

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Execution of this instruction is privileged.

The PPC440x5 does not automatically synchronize context between an **icread** instruction and the subsequent **mfspr** instructions which read the results of the **icread** instruction into GPRs. In order to guarantee that the **mfspr** instructions obtain the results of the **icread** instruction, a sequence such as the following must be used:

```

icread    regA,regB    # read cache information (the contents of GPR A and GPR B are
                    # added and the result used to specify a cache line index to be read)
isync
mficbdr  regC          # move instruction information into GPR C
mficbtrh regD          # move high portion of tag into GPR D
mficbtrl regE          # move low portion of tag into GPR E

```

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## isel


Integer Select

PPC440x5 CPU Core User's Manual



Preliminary

**isel**            RT, RA, RB, CRb



if CR[CRb] = 1 then

    (RT) ← (RA|0)

else

    (RT) ← (RB)

If CR[CRb] = 0, register RT is written with the contents of register RB.


If CR[CRb] = 1 and RA ≠ 0, register RT is written with the contents of register RA.

If CR[CRb] = 1 and RA = 0, register RT is written with 0.

### Registers Altered

- RT

## isync



The **isync** instruction is a context synchronizing instruction.

**isync** provides an ordering function for the effects of all instructions executed by the processor. Executing **isync** insures that all instructions preceding the **isync** instruction execute before **isync** completes, except that storage accesses caused by those instructions need not have completed. Furthermore, all instructions preceding the **isync** are guaranteed to be unaffected by any context changes initiated by instructions after the **isync**.

No subsequent instructions are initiated by the processor until **isync** completes. Finally, execution of **isync** causes the processor to discard any prefetched instructions (prefetched from the cache, not instructions that are in the cache or on their way into the cache), with the effect that subsequent instructions are fetched and executed in the context established by the instructions preceding **isync**.

**isync** causes any caching inhibited instruction fetches from memory to be aborted and any data associated with them to be discarded. Cacheable instruction fetches from memory are not aborted however, as these should be handled by the **icbi** instructions which must precede the **isync** if software wishes to invalidate any cached instructions.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Note

See the discussion of context synchronizing instructions in *Synchronization* on page 79.

The following code example illustrates the necessary steps for self-modifying code. This example assumes that `addr1` is both data and instruction cacheable.

```

    stw      regN, addr1          # data in regN is to become an instruction at addr1
    dcbst   addr1                # forces data from the data cache to memory
    msync                   # wait until the data actually reaches the memory
    icbi    addr1                # invalidate the instruction if it is in the cache (or in the #
process of being fetched into the cache)
    msync                   # wait until the icbi completes
    isync                   # discard and refetch any instructions (including
                           # possibly the instruction at addr1) which may have
                           # already been fetched from the cache and be in the
                           # pipeline after the isync

```

## lbz


Load Byte and Zero

PPC440x5 CPU Core User's Manual



Preliminary

**lbz**                    RT, D(RA)



$$EA \leftarrow (RA|0) + \text{EXTS}(D)$$
$$(RT) \leftarrow {}^{24}0 \parallel \text{MS}(EA,1)$$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.


The byte at the EA is extended to 32 bits by concatenating 24 0-bits to its left. The result is placed into register RT.

### Registers Altered

- RT



**lbzu**            RT, D(RA)



$$\begin{aligned}
 EA &\leftarrow (RA|0) + \text{EXTS}(D) \\
 (RA) &\leftarrow EA \\
 (RT) &\leftarrow {}^{24}0 \parallel \text{MS}(EA,1)
 \end{aligned}$$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The byte at the EA is extended to 32 bits by concatenating 24 0-bits to its left. The result is placed into register RT.

**Registers Altered**

- RA
- RT

**Invalid Instruction Forms**

- RA = RT
- RA = 0


## lbzux

Load Byte and Zero with Update Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lbzux** RT, RA, RB


$$\begin{aligned}EA &\leftarrow (RA|0) + (RB) \\(RA) &\leftarrow EA \\(RT) &\leftarrow 2^{24}0 \parallel MS(EA,1)\end{aligned}$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The byte at the EA is extended to 32 bits by concatenating 24 0-bits to its left. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered


- RA
- RT

### Invalid Instruction Forms

- Reserved fields
- RA = RT
- RA = 0



**lbzx** RT,RA, RB



$$EA \leftarrow (RA|0) + (RB)$$
$$(RT) \leftarrow {}^{24}0 \parallel MS(EA,1)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The byte at the EA is extended to 32 bits by concatenating 24 0-bits to its left. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RT

**Invalid Instruction Forms**

- Reserved fields


## lha

Load Halfword Algebraic  
PPC440x5 CPU Core User's Manual



Preliminary

**lha** RT, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $(RT) \leftarrow EXTS(MS(EA,2))$


An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The halfword at the EA is sign-extended to 32 bits and placed into register RT.

### Registers Altered

- RT

**lhau**            RT, D(RA)



EA ← (RA|0) + EXTS(D)  
 (RA) ← EA  
 (RT) ← EXTS(MS(EA,2))

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The halfword at the EA is sign-extended to 32 bits and placed into register RT.

**Registers Altered**

- RA
- RT

**Invalid Instruction Forms**

- RA = RT
- RA = 0


# lhaux

Load Halfword Algebraic with Update Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lhaux**            RT, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $(RA) \leftarrow EA$   
 $(RT) \leftarrow EXTS(MS(EA,2))$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The halfword at the EA is sign-extended to 32 bits and placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.


## Registers Altered

- RA
- RT

## Invalid Instruction Forms

- Reserved fields
- RA = RT
- RA = 0

**lhax**            RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The halfword at the EA is sign-extended to 32 bits and placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RT

**Invalid Instruction Forms**

- Reserved fields


# lhbrx

Load Halfword Byte-Reverse Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lhbrx** RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$
$$(RT) \leftarrow \overset{16}{0} \parallel \text{BYTE\_REVERSE}(\text{MS}(EA,2))$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The halfword at the EA is byte-reversed from the default byte ordering for the memory page referenced by the EA. The resulting halfword is extended to 32 bits by concatenating 16 0-bits to its left. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT

### Invalid Instruction Forms

- Reserved fields


### Programming Note

Byte ordering is generally controlled by the Endian (E) storage attribute (see *Memory Management* on page 129). The load byte reverse instructions provide a mechanism for data to be loaded from a memory page using the opposite byte ordering from that specified by the Endian storage attribute.





**lhz** RT, D(RA)



$$EA \leftarrow (RA|0) + EXTS(D)$$
$$(RT) \leftarrow {}^{16}0 \parallel MS(EA,2)$$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The halfword at the EA is extended to 32 bits by concatenating 16 0-bits to its left. The result is placed into register RT.

**Registers Altered**

- RT


# lhzu

Load Halfword and Zero with Update  
PPC440x5 CPU Core User's Manual



Preliminary

**lhzu** RT, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $(RA) \leftarrow EA$   
 $(RT) \leftarrow {}^{16}0 \parallel MS(EA,2)$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.


The halfword at the EA is extended to 32 bits by concatenating 16 0-bits to its left. The result is placed into register RT.

## Registers Altered

- RA
- RT

## Invalid Instruction Forms

- RA = RT
- RA = 0

**lhzux RT, RA, RB**


$$EA \leftarrow (RA|0) + (RB)$$

$$(RA) \leftarrow EA$$

$$(RT) \leftarrow {}^{16}0 \parallel MS(EA,2)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The halfword at the EA is extended to 32 bits by concatenating 16 0-bits to its left. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RA
- RT

**Invalid Instruction Forms**

- Reserved fields
- RA = RT
- RA = 0


## lhzx

Load Halfword and Zero Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lhzx**            RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$
$$(RT) \leftarrow {}^{16}0 \parallel MS(EA,2)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The halfword at the EA is extended to 32 bits by concatenating 16 0-bits to its left. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.


### Registers Altered

- RT

### Invalid Instruction Forms

- Reserved fields

**Imw**            RT, D(RA)



```
EA ← (RA|0) + EXTS(D)
r ← RT
do while r ≤ 31
  GPR(r) ← MS(EA,4)
  r ← r + 1
  EA ← EA + 4
```

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field in the instruction to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

A series of consecutive words starting at the EA are loaded into a set of consecutive GPRs, starting with register RT and continuing to and including GPR(31).

**Registers Altered**

- RT through GPR(31).

**Invalid Instruction Forms**

- RA is in the range of registers to be loaded, including the case RA = RT = 0.

**Programming Note**

This instruction can be restarted, meaning that it could be interrupted after having already updated some of the target registers, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been loaded prior to the interrupt will be loaded a second time. Note that if RA is in the range of registers to be loaded (an invalid form; see above) and is also one of the registers which is loaded prior to the interrupt, then when the instruction is restarted the re-calculated EA will be incorrect, since RA will no longer contain the original base address. Hence the definition of this as an invalid form which software must avoid.


## lswi

Load String Word Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

**lswi**            RT, RA, NB



```
EA ← (RA|0)
if NB = 0 then
    CNT ← 32
else
    CNT ← NB
n ← CNT
RFINAL ← ((RT + CEIL(CNT/4) - 1) % 32)
r ← RT - 1
i ← 0
do while n > 0
    if i = 0 then
        r ← r + 1
        if r = 32 then
            r ← 0
        GPR(r) ← 0
    GPR(r)i:i+7 ← MS(EA,1)
    i ← i + 8
    if i = 32 then
        i ← 0
    EA ← EA + 1
    n ← n - 1
```

An effective address (EA) is determined by the RA field. If the RA field contains 0, the EA is 0. Otherwise, the EA is the contents of register RA.

The NB field specifies the byte count CNT. If the NB field contains 0, the byte count is CNT = 32. Otherwise, the byte count is CNT = NB.

A series of CNT consecutive bytes in main storage, starting at the EA, are loaded into CEIL(CNT/4) consecutive GPRs, four bytes per GPR, until the byte count is exhausted. Bytes are loaded into GPRs; the byte at the lowest address is loaded into the most significant byte. Bits to the right of the last byte loaded into the last GPR are set to 0.

The set of loaded GPRs starts at register RT, continues consecutively through GPR(31), and wraps to register 0, loading until the byte count is exhausted, which occurs in register R<sub>FINAL</sub>.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT and subsequent GPRs as described above.

### Invalid Instruction Forms

- Reserved fields
- RA is in the range of registers to be loaded

**Preliminary**

---

- RA = RT = 0

**Programming Note**

This instruction can be restarted, meaning that it could be interrupted after having already updated some of the target registers, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been loaded prior to the interrupt will be loaded a second time. Note that if RA is in the range of registers to be loaded (an invalid form; see above) and is also one of the registers which is loaded prior to the interrupt, then when the instruction is restarted the re-calculated EA will be incorrect, since RA will no longer contain the original base address. Hence the definition of this as an invalid form which software must avoid.


## lswx

Load String Word Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lswx** RT, RA, RB



```
EA ← (RA|0) + (RB)
CNT ← XER[TBC]
n ← CNT
RFINAL ← ((RT + CEIL(CNT/4) – 1) % 32)
r ← RT – 1
i ← 0
do while n > 0
  if i = 0 then
    r ← r + 1
    if r = 32 then
      r ← 0
    GPR(r) ← 0
  GPR(r)i:i+7 ← MS(EA,1)
  i ← i + 8
  if i = 32 then
    i ← 0
  EA ← EA + 1
  n ← n – 1
```

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

A byte count CNT is obtained from XER[TBC].

A series of CNT consecutive bytes in main storage, starting at the EA, are loaded into CEIL(CNT/4) consecutive GPRs, four bytes per GPR, until the byte count is exhausted. Bytes are loaded into GPRs; the byte having the lowest address is loaded into the most significant byte. Bits to the right of the last byte loaded in the last GPR used are set to 0.

The set of consecutive GPRs loaded starts at register RT, continues through GPR(31), and wraps to register 0, loading until the byte count is exhausted, which occurs in register R<sub>FINAL</sub>.

If XER[TBC] is 0, the byte count is 0 and the contents of register RT are undefined.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT and subsequent GPRs as described above.

### Invalid Instruction Forms

- Reserved fields
- RA or RB is in the range of registers to be loaded.
- RA = RT = 0



**Programming Note**

This instruction can be restarted, meaning that it could be interrupted after having already updated some of the target registers, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been loaded prior to the interrupt will be loaded a second time. Note that if RA or RB is in the range of registers to be loaded (an invalid form; see above) and is also one of the registers which is loaded prior to the interrupt, then when the instruction is restarted the re-calculated EA will be incorrect, since the affected register will no longer contain the original base address or index. Hence the definition of these as invalid forms which software must avoid.

If XER[TBC] = 0, the contents of register RT are undefined and **lswx** is treated as a no-op. Furthermore, if the EA is such that a Data Storage, Data TLB Error, or Data Address Compare Debug exception occurs, **lswx** is treated as a no-op and no interrupt occurs as a result of the exception.


# lwarx

Load Word and Reserve Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lwarx** RT, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $RESERVE \leftarrow 1$   
 $(RT) \leftarrow MS(EA,4)$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The word at the EA is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Execution of the **lwarx** instruction sets the reservation bit.

## Registers Altered

- RT

## Invalid Instruction Forms

- Reserved fields


## Programming Note

The **lwarx** and **stwcx.** instructions are typically paired in a loop, as shown in the following example, to create the effect of an atomic operation to a memory area used as a semaphore between multiple processes. Only **lwarx** can set the reservation bit to 1. **stwcx.** sets the reservation bit to 0 upon its completion, whether or not **stwcx.** actually stored (RS) to memory. CR[CR0]<sub>2</sub> must be examined to determine whether (RS) was sent to memory.

```
loop: lwarx      # read the semaphore from memory; set reservation
      "alter"   # change the semaphore bits in the register as required
      stwcx.    # attempt to store the semaphore; reset reservation
      bne loop  # some other process intervened and cleared the reservation prior to the above
                        # stwcx.; try again
```

The PowerPC Book-E architecture specifies that the EA for the **lwarx** instruction must be word-aligned (that is, a multiple of 4 bytes); otherwise, the result is undefined. Although the PPC440x5 will execute **lwarx** regardless of the EA alignment, in order for the operation of the pairing of **lwarx** and **stwcx.** to produce the desired result, software must ensure that the EA for both instructions is word-aligned. This requirement is due to the manner in which misaligned storage accesses may be broken up into separate, aligned accesses by the PPC440x5.

**lwbrx**            RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$(RT) \leftarrow \text{BYTE\_REVERSE}(\text{MS}(EA,4))$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The word at the EA is byte-reversed from the default byte ordering for the memory page referenced by the EA. The result is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RT

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Byte ordering is generally controlled by the Endian (E) storage attribute (see *Memory Management* on page 129). The load byte reverse instructions provide a mechanism for data to be loaded from a memory page using the opposite byte ordering from that specified by the Endian storage attribute.

## lwz


Load Word and Zero

PPC440x5 CPU Core User's Manual



Preliminary

**lwz** RT, D(RA)



$EA \leftarrow (RA|0) + \text{EXTS}(D)$   
 $(RT) \leftarrow \text{MS}(EA,4)$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.


The word at the EA is placed into register RT.

### Registers Altered

- RT



**lwzu** RT, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $(RA) \leftarrow EA$   
 $(RT) \leftarrow MS(EA,4)$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The word at the EA is placed into register RT.

**Registers Altered**

- RA
- RT

**Invalid Instruction Forms**

- RA = RT
- RA = 0


## lwzux

Load Word and Zero with Update Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**lwzux** RT, RA, RB



$EA \leftarrow (RA|0) + (RB)$

$(RA) \leftarrow EA$

$(RT) \leftarrow MS(EA,4)$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise. The EA is placed into register RA.

The word at the EA is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.


### Registers Altered

- RA
- RT

### Invalid Instruction Forms

- Reserved fields
- RA = RT
- RA = 0

**lwzx**            RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$(RT) \leftarrow MS(EA,4)$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The word at the EA is placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RT

**Invalid Instruction Forms**

- Reserved fields

## macchw

Multiply Accumulate Cross Halfword to Word Modulo Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>macchw</b>	RT, RA, RB	OE=0, Rc=0
<b>macchw.</b>	RT, RA, RB	OE=0, Rc=1
<b>macchwo</b>	RT, RA, RB	OE=1, Rc=0
<b>macchwo.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	172	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$  signed

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is summed with the contents of RT and RT is updated with the low-order 32 bits of the signed sum.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.





Preliminary

<b>macchws</b>	RT, RA, RB	OE=0, Rc=0
<b>macchws.</b>	RT, RA, RB	OE=0, Rc=1
<b>macchwso</b>	RT, RA, RB	OE=1, Rc=0
<b>macchwso.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	236	Rc
0	6	11	16	21 22		31

```

prod0:31 ← (RA)16:31 × (RB)0:15 signed
temp0:32 ← prod0:31 + (RT)
if ((prod0 = RT0) ∧ (RT0 ≠ temp1)) then (RT) ← (RT0 || 31(¬RT0))
else (RT) ← temp1:32

```

The low-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is summed with the contents of RT.

If the signed sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the signed sum.

If the signed sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the signed sum is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the signed sum is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## macchwsu

Multiply Accumulate Cross Halfword to Word Saturate Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

<b>macchwsu</b>	RT, RA, RB	OE=0, Rc=0
<b>macchwsu.</b>	RT, RA, RB	OE=0, Rc=1
<b>macchwsuo</b>	RT, RA, RB	OE=1, Rc=0
<b>macchwsuo.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	204	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$  unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow (temp_{1:32} \vee^{32} temp_0)$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The unsigned product is summed with the contents of RT.

If the unsigned sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the unsigned sum.

If the unsigned sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the maximum representable value of  $2^{32} - 1$ .

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



Preliminary

<b>macchwu</b>	RT, RA, RB	OE=0, Rc=0
<b>macchwu.</b>	RT, RA, RB	OE=0, Rc=1
<b>macchwuo</b>	RT, RA, RB	OE=1, Rc=0
<b>macchwuo.</b>	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$  unsigned  
 $temp_{0:32} \leftarrow prod_{0:31} + (RT)$   
 $(RT) \leftarrow temp_{1:32}$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The unsigned product is summed with the contents of RT and RT is updated with the low-order 32 bits of the unsigned sum.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## **machhw**

Multiply Accumulate High Halfword to Word Modulo Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>machhw</b>	RT, RA, RB	OE=0, Rc=0
<b>machhw.</b>	RT, RA, RB	OE=0, Rc=1
<b>machhwo</b>	RT, RA, RB	OE=1, Rc=0
<b>machhwo.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	44	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$  signed

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is summed with the contents of RT and RT is updated with the low-order 32 bits of the signed sum.

### **Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

### **Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>machhws</b>	RT, RA, RB	OE=0, Rc=0
<b>machhws.</b>	RT, RA, RB	OE=0, Rc=1
<b>machhwso</b>	RT, RA, RB	OE=1, Rc=0
<b>machhwso.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	108	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$  signed  
 $temp_{0:32} \leftarrow prod_{0:31} + (RT)$   
 if  $((prod_0 = RT_0) \wedge (RT_0 \neq temp_1))$  then  $(RT) \leftarrow (RT_0 \parallel^{31}(\neg RT_0))$   
 else  $(RT) \leftarrow temp_{1:32}$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is summed with the contents of RT.

If the signed sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the signed sum.

If the signed sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the signed sum is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the signed sum is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## **machhwsu**

Multiply Accumulate High Halfword to Word Saturate Unsigned  
**PPC440x5 CPU Core User's Manual**



Preliminary

<b>machhwsu</b>	RT, RA, RB	OE=0, Rc=0
<b>machhwsu.</b>	RT, RA, RB	OE=0, Rc=1
<b>machhwsuo</b>	RT, RA, RB	OE=1, Rc=0
<b>machhwsuo.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	76	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$  unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow (temp_{1:32} \vee^{32} temp_0)$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The unsigned product is summed with the contents of RT.

If the unsigned sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the unsigned sum.

If the unsigned sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the maximum representable value of  $2^{32} - 1$ .

### **Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### **Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>machhwu</b>	RT, RA, RB	OE=0, Rc=0
<b>machhwu.</b>	RT, RA, RB	OE=0, Rc=1
<b>machhwuo</b>	RT, RA, RB	OE=1, Rc=0
<b>machhwuo.</b>	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$  unsigned  
 $temp_{0:32} \leftarrow prod_{0:31} + (RT)$   
 $(RT) \leftarrow temp_{1:32}$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The unsigned product is summed with the contents of RT and RT is updated with the low-order 32 bits of the unsigned sum.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## macihw

Multiply Accumulate Low Halfword to Word Modulo Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>macihw</b>	RT, RA, RB	OE=0, Rc=0
<b>macihw.</b>	RT, RA, RB	OE=0, Rc=1
<b>maciwho</b>	RT, RA, RB	OE=1, Rc=0
<b>maciwho.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	428	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$  signed

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$

The low-order halfword of RA is multiplied by the low-order halfword of RB. The signed product is summed with the contents of RT and RT is updated with the low-order 32 bits of the signed sum.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>maclhws</b>	RT, RA, RB	OE=0, Rc=0
<b>maclhws.</b>	RT, RA, RB	OE=0, Rc=1
<b>maclhwso</b>	RT, RA, RB	OE=1, Rc=0
<b>maclhwso.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	492	Rc
0	6	11	16	21 22		31

```

prod0:31 ← (RA)16:31 × (RB)16:31 signed
temp0:32 ← prod0:31 + (RT)
if ((prod0 = RT0) ∧ (RT0 ≠ temp1)) then (RT) ← (RT0 || 31(¬RT0))
else (RT) ← temp1:32
    
```

The low-order halfword of RA is multiplied by the low-order halfword of RB. The signed product is summed with the contents of RT.

If the signed sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the signed sum.

If the signed sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the signed sum is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the signed sum is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## maclhwsu

Multiply Accumulate Low Halfword to Word Saturate Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

<b>maclhwsu</b>	RT, RA, RB	OE=0, Rc=0
<b>maclhwsu.</b>	RT, RA, RB	OE=0, Rc=1
<b>maclhwsuo</b>	RT, RA, RB	OE=1, Rc=0
<b>maclhwsuo.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	460	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$  unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow (temp_{1:32} \vee^{32} temp_0)$

The low-order halfword of RA is multiplied by the low-order halfword of RB. The unsigned product is summed with the contents of RT.

If the unsigned sum can be represented in 32 bits, then RT is updated with the low-order 32 bits of the unsigned sum.

If the unsigned sum cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the maximum representable value of  $2^{32} - 1$ .

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>maclhwu</b>	RT, RA, RB	OE=0, Rc=0
<b>maclhwu.</b>	RT, RA, RB	OE=0, Rc=1
<b>maclhwuo</b>	RT, RA, RB	OE=1, Rc=0
<b>maclhwuo.</b>	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$  unsigned  
 $temp_{0:32} \leftarrow prod_{0:31} + (RT)$   
 $(RT) \leftarrow temp_{1:32}$

The low-order halfword of RA is multiplied by the low-order halfword of RB. The unsigned product is summed with the contents of RT and RT is updated with the low-order 32 bits of the unsigned sum.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

# mbar



## mbar



The **mbar** instruction ensures that all loads and stores preceding **mbar** complete with respect to main storage before any loads and stores following **mbar** access main storage. As implemented in the PPC440x5 core, the MO field of **mbar** is ignored and treated as 0, providing a storage ordering function for all storage access instructions executed by the processor. Other processors implementing the **mbar** instruction may support one or more non-zero MO settings, specifying different subsets of storage accesses to be ordered by the **mbar** instruction in those implementations.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

### Programming Note

Architecturally, **mbar** merely orders storage accesses, and does not perform execution nor context synchronization (see *Synchronization* on page 79). Therefore, non-storage access instructions after **mbar** could complete before the storage access instructions which were executed prior to **mbar** have actually completed their storage accesses. The **msync** instruction, on the other hand, *is* execution synchronizing, and *does* guarantee that all storage accesses initiated by instructions executed prior to the **msync** have completed before any instructions after the **msync** begin execution. However, the PPC440x5 core implements the **mbar** instruction identically to the **msync** instruction, and thus both are execution synchronizing.

Software should nevertheless use the correct instruction (**mbar** or **msync**) as called for by the specific ordering and synchronizing requirements of the application, in order to guarantee portability to other implementations.

See *Storage Ordering and Synchronization* on page 81 for additional information on the use of the **msync** and **mbar** instructions.

Table 9-19. Extended Mnemonics for mbar

Mnemonic	Operands	Function	Other Registers Altered
<b>mbar</b>	None	Memory Barrier. <i>Extended mnemonic for mbar 0</i>	



---

**Architecture Note**

**mbar** replaces the PowerPC **eiio** instruction. **mbar** uses the same opcode as **eiio**; PowerPC applications which used **eiio** will get the function of **mbar** when executed on a PowerPC Book-E implementation. **mbar** is architecturally “stronger” than **eiio**, in that **eiio** forced separate ordering amongst different *categories* of storage accesses, while **mbar** forces such ordering amongst *all* storage accesses as a single category.


## mcrf

Move Condition Register Field  
PPC440x5 CPU Core User's Manual



Preliminary

**mcrf**      BF, BFA



$m \leftarrow \text{BFA}$   
 $n \leftarrow \text{BF}$   
 $(\text{CR}[\text{CR}_n]) \leftarrow (\text{CR}[\text{CR}_m])$

The contents of the CR field specified by the BFA field are placed into the CR field specified by the BF field.

### Registers Altered


- $\text{CR}[\text{CR}_n]$  where  $n$  is specified by the BF field.

### Invalid Instruction Forms

- Reserved fields



**mcrxr**      BF



$n \leftarrow \text{BF}$   
 $(\text{CR}[\text{CR}_n]) \leftarrow \text{XER}_{0:3}$   
 $\text{XER}_{0:3} \leftarrow 40$

The contents of  $\text{XER}_{0:3}$  are placed into the CR field specified by the BF field.  $\text{XER}_{0:3}$  are then set to 0.

If instruction bit 31 contains 1, the contents of  $\text{CR}[\text{CR}_0]$  are undefined.

**Registers Altered**

- $\text{CR}[\text{CR}_n]$  where  $n$  is specified by the BF field.
- $\text{XER}[\text{SO}, \text{OV}, \text{CA}]$

**Invalid Instruction Forms**

- Reserved fields


## mfcrr

Move From Condition Register  
PPC440x5 CPU Core User's Manual



Preliminary

**mfcrr**            RT



$(RT) \leftarrow (CR)$

The contents of the CR are placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT

### Invalid Instruction Forms

- Reserved fields



**mfdcr**      RT, DCRN


$$\text{DCRN} \leftarrow \text{DCRF}_{5:9} \parallel \text{DCRF}_{0:4}$$

$$(\text{RT}) \leftarrow (\text{DCR}(\text{DCRN}))$$

The contents of the DCR specified by the DCRF field are placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

Execution of this instruction is privileged.

The DCR number (DCRN) specified in the assembler language coding of the **mfdcr** instruction refers to a DCR number. The assembler handles the unusual register number encoding to generate the DCRF field.

### Architecture Note

The specific numbers and definitions of any DCRs are outside the scope of both the PowerPC Book-E architecture and the PPC440x5 core. Any DCRs are defined as part of the chip-level product incorporating the PPC440x5 core.


## mfmsr

Move From Machine State Register  
PPC440x5 CPU Core User's Manual



Preliminary

**mfmsr**      RT



$(RT) \leftarrow (MSR)$

The contents of the MSR are placed into register RT.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- RT

### Invalid Instruction Forms

- Reserved fields

### Programming Note

Execution of this instruction is privileged.

**mfspr**      RT, SPRN



$$\begin{aligned} \text{SPRN} &\leftarrow \text{SPRF}_{5:9} \parallel \text{SPRF}_{0:4} \\ (\text{RT}) &\leftarrow (\text{SPR}(\text{SPRN})) \end{aligned}$$

The contents of the SPR specified by the SPRF field are placed into register RT. See *Special Purpose Registers Sorted by SPR Number* on page 443 for a listing of SPR mnemonics and corresponding SPRN values.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RT

**Invalid Instruction Forms**

- Reserved fields
- Invalid SPRF values

**Programming Note**

Execution of this instruction is privileged if instruction bit 11 contains 1. See *Privileged SPRs* on page 79 for a list of privileged SPRs.

The SPR number (SPRN) specified in the assembler language coding of the **mfspr** instruction refers to an SPR number. The assembler handles the unusual register number encoding to generate the SPRF field.

# mfspr

Move From Special Purpose Register  
 PPC440x5 CPU Core User's Manual



Preliminary

Table 9-20. Extended Mnemonics for *mfspr*

Mnemonic	Operands	Function
mfccr0 mfccr1 mfcsrr0 mfcsrr1 mfctr mfdac1 mfdac2 mfdbcr0 mfdbcr1 mfdbcr2 mfdbdr mfdbsr mfdcdbtrh mfdcdbtrl mfdear mfdec mfdnv0 mfdnv1 mfdnv2 mfdnv3 mfdtv0 mfdtv1 mfdtv2 mfdtv3 mfdivc1 mfdivc2 mfdivlim mfesr mfia1 mfia2 mfia3 mfia4 mficbdr mficbtrh mficbtrl mfinv0 mfinv1 mfinv2 mfinv3 mfity0 mfity1 mfity2 mfity3 mfivlim mfivor0 mfivor1 mfivor2 mfivor3 mfivor4 mfivor5 mfivor6 mfivor7 mfivor8 mfivor9 mfivor10 mfivor11 mfivor12 mfivor13 mfivor14 mfivor15 mfivpr mflr mfmcscr mfmcscr0 mfmcscr1 mfmmucr	RT	<p>Move from special purpose register SPRN.  <i>Extended mnemonic for</i>  <b>mfspr RT,SPRN</b></p> <p>See <i>Special Purpose Registers Sorted by SPR Number</i> on page 443 for a list of valid SPRN values.</p>



Table 9-20. Extended Mnemonics for mfspr (continued)

Mnemonic	Operands	Function
<b>mfpid</b> <b>mfpir</b> <b>mfpvr</b> <b>mfsprg0</b> <b>mfsprg1</b> <b>mfsprg2</b> <b>mfsprg3</b> <b>mfsprg4</b> <b>mfsprg5</b> <b>mfsprg6</b> <b>mfsprg7</b> <b>mfsrr0</b> <b>mfsrr1</b> <b>mftbl</b> <b>mftbu</b> <b>mftcr</b> <b>mftsr</b> <b>mfusprg0</b> <b>mfxer</b>		

# msync


Memory Synchronize

PPC440x5 CPU Core User's Manual



Preliminary

## msync



The **msync** instruction guarantees that all instructions initiated by the processor preceding **msync** will complete before **msync** completes, and that no subsequent instructions will be initiated by the processor until after **msync** completes. **msync** also will not complete until all storage accesses associated with instructions preceding **msync** have completed.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None.

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

The **msync** instruction is execution synchronizing (see *Execution Synchronization* on page 81), and guarantees that all storage accesses initiated by instructions executed prior to the **msync** have completed before any instructions after the **msync** begin execution. On the other hand, architecturally the **mbar** instruction merely *orders* storage accesses, and does *not* perform execution synchronization. Therefore, non-storage access instructions after **mbar** *could* complete before the storage access instructions which were executed prior to **mbar** have actually completed their storage accesses. However, the PPC440x5 core implements the **mbar** instruction identically to the **msync** instruction, and thus both are execution synchronizing.


Software should nevertheless use the correct instruction (**mbar** or **msync**) as called for by the specific ordering and synchronizing requirements of the application, in order to guarantee portability to other implementations.

See *Storage Ordering and Synchronization* on page 81 for additional information on the use of the **msync** and **mbar** instructions.

### Architecture Note

**mbar** replaces the PowerPC **eieio** instruction. **mbar** uses the same opcode as **eieio**; PowerPC applications which used **eieio** will get the function of **mbar** when executed on a PowerPC Book-E implementation. **mbar** is architecturally “stronger” than **eieio**, in that **eieio** forced separate ordering amongst different *categories* of storage accesses, while **mbar** forces such ordering amongst *all* storage accesses as a single category.

**msync** replaces the PowerPC **sync** instruction. **msync** uses the same opcode as **sync**; PowerPC applications which used **sync** get the function of **msync** when executed on a PowerPC Book-E implementation. **msync** is architecturally identical to the version of **sync** specified by an earlier version of the PowerPC architecture.

**mtcrf**      FXM, RS


$$\text{mask} \leftarrow {}^4(\text{FXM}_0) \parallel {}^4(\text{FXM}_1) \parallel \dots \parallel {}^4(\text{FXM}_6) \parallel {}^4(\text{FXM}_7)$$

$$(\text{CR}) \leftarrow ((\text{RS}) \wedge \text{mask}) \vee ((\text{CR}) \wedge \neg \text{mask})$$

Some or all of the contents of register RS are placed into the CR as specified by the FXM field.

Each bit in the FXM field controls the copying of 4 bits in register RS into the corresponding bits in the CR. The correspondence between the bits in the FXM field and the bit copying operation is shown in the following table:

*Table 9-21. FXM Bit Field Correspondence*

FXM Bit Number	CR Bits Affected
0	0:3
1	4:7
2	8:11
3	12:15
4	16:19
5	20:23
6	24:27
7	28:31

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- CR

### Invalid Instruction Forms

- Reserved fields

*Table 9-22. Extended Mnemonics for mtcrf*

Mnemonic	Operands	Function
<b>mtcr</b>	RS	Move to CR. <i>Extended mnemonic for</i> <b>mtcrf 0xFF,RS</b>

## mtdcr

Move To Device Control Register  
PPC440x5 CPU Core User's Manual



Preliminary

**mtdcr**          DCRN, RS



$DCRN \leftarrow DCRF_{5:9} \parallel DCRF_{0:4}$   
 $(DCR(DCRN)) \leftarrow (RS)$

The contents of register RS are placed into the DCR specified by the DCRF field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- DCR(DCRN)

### Invalid Instruction Forms

- Reserved fields

### Programming Note

Execution of this instruction is privileged.

The DCR number (DCRN) specified in the assembler language coding of the **mtdcr** instruction refers to a DCR number. The assembler handles the unusual register number encoding to generate the DCRF field.


### Architecture Note

The specific numbers and definitions of any DCRs are outside the scope of both the PowerPC Book-E architecture and the PPC440x5 core. Any DCRs are defined as part of the chip-level product incorporating the PPC440x5 core.





**mtmsr**      RS



(MSR) ← (RS)

The contents of register RS are placed into the MSR.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- MSR

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

The **mtmsr** instruction is privileged and execution synchronizing (see *Execution Synchronization* on page 81).

## mtspr

Move To Special Purpose Register  
PPC440x5 CPU Core User's Manual



Preliminary

**mtspr**          SPRN, RS


$$\text{SPRN} \leftarrow \text{SPRF}_{5:9} \parallel \text{SPRF}_{0:4}$$
$$(\text{SPR}(\text{SPRN})) \leftarrow (\text{RS})$$

The contents of register RS are placed into register RT. See *Special Purpose Registers Sorted by SPR Number* on page 443 for a listing of SPR mnemonics and corresponding SPRN values.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- SPR (SPRN)

### Invalid Instruction Forms

- Reserved fields
- Invalid SPRF values

### Programming Note

Execution of this instruction is privileged if instruction bit 11 contains 1. See *Privileged SPRs* on page 79 for a list of privileged SPRs.

The SPR number (SPRN) specified in the assembler language coding of the **mtspr** instruction refers to an SPR number. The assembler handles the unusual register number encoding to generate the SPRF field.

Table 9-23. Extended Mnemonics for mtspr

Mnemonic	Operands	Function
mtccr0 mtccr1 mtcsrr0 mtcsrr1 mtctr mtdac1 mtdac2 mtdbcr0 mtdbcr1 mtdbcr2 mtdbdr mtdbsr mtdear mtdec mtdecar mtdnv0 mtdnv1 mtdnv2 mtdnv3 mtdtv0 mtdtv1 mtdtv2 mtdtv3 mtdvc1 mtdvc2 mtdvlim mtesr mtiac1 mtiac2 mtiac3 mtiac4 mtinv0 mtinv1 mtinv2 mtinv3 mtitv0 mtitv1 mtitv2 mtitv3 mtivlim mtivor0 mtivor1 mtivor2 mtivor3 mtivor4 mtivor5 mtivor6 mtivor7 mtivor8 mtivor9 mtivor10 mtivor11 mtivor12 mtivor13 mtivor14 mtivor15 mtivpr mtlr mtmcsr mtmcsrr0 mtmcsrr1 mtmmucr mtpid	RT	<p>Move to special purpose register SPRN. <i>Extended mnemonic for</i> <b>mtspr RT,SPRN</b></p> <p>See <i>Special Purpose Registers Sorted by SPR Number</i> on page 443 for a list of valid SPRN values.</p>

# mtspr

Move To Special Purpose Register  
PPC440x5 CPU Core User's Manual




Preliminary

Table 9-23. Extended Mnemonics for mtspr (continued)

Mnemonic	Operands	Function
<b>mtsprg0</b> <b>mtsprg1</b> <b>mtsprg2</b> <b>mtsprg3</b> <b>mtsprg4</b> <b>mtsprg5</b> <b>mtsprg6</b> <b>mtsprg7</b> <b>mtsrr0</b> <b>mtsrr1</b> <b>mttbl</b> <b>mttbu</b> <b>mttcr</b> <b>mttsr</b> <b>mtusprg0</b> <b>mtxer</b>		



**mulchw** RT, RA, RB Rc=0  
**mulchw.** RT, RA, RB Rc=1



$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15} \text{ signed}$$

The low-order halfword of RA is multiplied by the high-order halfword of RB, considering both source operands as signed integers. The 32-bit result is placed into register RT.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.


## mulchwu

Multiply Cross Halfword to Word Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

**mulchwu**      RT, RA, RB                      Rc=0  
**mulchwu.**     RT, RA, RB                      Rc=1



$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15} \text{ unsigned}$$

The low-order halfword of RA is multiplied by the high-order halfword of RB, considering both source operands as unsigned integers. The 32-bit result is placed into register RT.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1


### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

**mulhww**      RT, RA, RB                      Rc=0  
**mulhww.**     RT, RA, RB                      Rc=1



$$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15} \text{ signed}$$

The high-order halfword of RA is multiplied by the high-order halfword of RB, considering both source operands as signed integers. The 32-bit result is placed into register RT.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.


## mulhhwu

Multiply High Halfword to Word Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

<b>mulhhwu</b>	RT, RA, RB	Rc=0
<b>mulhhwu.</b>	RT, RA, RB	Rc=1



$$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15} \text{ unsigned}$$

The high-order halfword of RA is multiplied by the high-order halfword of RB, considering both source operands as unsigned integers. The 32-bit result is placed into register RT.

### Registers Altered


- RT
- CR[CR0] if Rc contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



<b>mulhw</b>	RT, RA, RB	Rc=0
<b>mulhw.</b>	RT, RA, RB	Rc=1



$prod_{0:63} \leftarrow (RA) \times (RB) \text{ signed}$   
 $(RT) \leftarrow prod_{0:31}$

The 64-bit signed product of registers RA and RB is formed. The most significant 32 bits of the result is placed into register RT.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1

**Programming Note**

The most significant 32 bits of the product, unlike the least significant 32 bits, may differ depending on whether the registers RA and RB are interpreted as signed or unsigned quantities. **mulhw** generates the correct result when these operands are interpreted as signed quantities. **mulhwu** generates the correct result when these operands are interpreted as unsigned quantities.

**Invalid Instruction Forms**

- Reserved fields


## mulhwu

Multiply High Word Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

<b>mulhwu</b>	RT, RA, RB	Rc=0
<b>mulhwu.</b>	RT, RA, RB	Rc=1



$prod_{0:63} \leftarrow (RA) \times (RB)$  unsigned  
 $(RT) \leftarrow prod_{0:31}$

The 64-bit unsigned product of registers RA and RB is formed. The most significant 32 bits of the result are placed into register RT.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1

### Programming Note

The most significant 32 bits of the product, unlike the least significant 32 bits, may differ depending on whether the registers RA and RB are interpreted as signed or unsigned quantities. The **mulhw** instruction generates the correct result when these operands are interpreted as signed quantities. The **mulhwu** instruction generates the correct result when these operands are interpreted as unsigned quantities.


### Invalid Instruction Forms

- Reserved fields



Preliminary

<b>mullhw</b>	RT, RA, RB	Rc=0
<b>mullhw.</b>	RT, RA, RB	Rc=1



$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31} \text{ signed}$$

The low-order halfword of RA is multiplied by the low-order halfword of RB, considering both source operands as signed integers. The 32-bit result is placed into register RT.

#### Registers Altered

- RT
- CR[CR0] if Rc contains 1

#### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.


## mullhwu

Multiply Low Halfword to Word Unsigned  
PPC440x5 CPU Core User's Manual



Preliminary

**mullhwu** RT, RA, RB R<sub>c</sub>=0  
**mullhwu.** RT, RA, RB R<sub>c</sub>=1



$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31} \text{ unsigned}$$

The low-order halfword of RA is multiplied by the low-order halfword of RB, considering both source operands as unsigned integers. The 32-bit result is placed into register RT.


### Registers Altered

- RT
- CR[CR0] if R<sub>c</sub> contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

**mulli**            RT, RA, IM



$$\begin{aligned} \text{prod}_{0:47} &\leftarrow (\text{RA}) \times \text{IM} \\ (\text{RT}) &\leftarrow \text{prod}_{16:47} \end{aligned}$$

The 48-bit product of register RA and the 16-bit IM field is formed. The least significant 32 bits of the product are placed into register RT.

**Registers Altered**

- RT

**Programming Note**

The least significant 32 bits of the product are correct, regardless of whether register RA and field IM are interpreted as signed or unsigned numbers.

## mullw


Multiply Low Word

PPC440x5 CPU Core User's Manual



Preliminary

<b>mullw</b>	RT, RA, RB	OE=0, Rc=0
<b>mullw.</b>	RT, RA, RB	OE=0, Rc=1
<b>mullwo</b>	RT, RA, RB	OE=1, Rc=0
<b>mullwo.</b>	RT, RA, RB	OE=1, Rc=1



$prod_{0:63} \leftarrow (RA) \times (RB) \text{ signed}$   
 $(RT) \leftarrow prod_{32:63}$

The 64-bit signed product of register RA and register RB is formed. The least significant 32 bits of the result is placed into register RT.

If the signed product cannot be represented in 32 bits and OE=1, XER[SO, OV] are set to 1.

### Registers Altered


- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE=1

### Programming Note

The least significant 32 bits of the product are correct, regardless of whether register RA and register RB are interpreted as signed or unsigned numbers. The overflow indication, however, is calculated specifically for a 64-bit signed product, and is dependent upon interpretation of the source operands as signed numbers.



**nand** RA, RS, RB Rc=0  
**nand.** RA, RS, RB Rc=1



$$(RA) \leftarrow \neg((RS) \wedge (RB))$$

The contents of register RS is ANDed with the contents of register RB; the ones complement of the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if Rc contains 1

# neg


Negate

PPC440x5 CPU Core User's Manual



Preliminary

<b>neg</b>	RT, RA	OE=0, Rc=0
<b>neg.</b>	RT, RA	OE=0, Rc=1
<b>nego</b>	RT, RA	OE=1, Rc=0
<b>nego.</b>	RT, RA	OE=1, Rc=1



$$(RT) \leftarrow \neg(RA) + 1$$

The two's complement of the contents of register RA are placed into register RT.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE=1

### Invalid Instruction Forms


- Reserved fields





**Preliminary**

<b>nmacchw</b>	RT, RA, RB	OE=0, Rc=0
<b>nmacchw.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmacchw0</b>	RT, RA, RB	OE=1, Rc=0
<b>nmacchw0.</b>	RT, RA, RB	OE=1, Rc=1



$$\begin{aligned} \text{nprod}_{0:31} &\leftarrow -((\text{RA})_{16:31} \times (\text{RB})_{0:15}) \text{ signed} \\ \text{temp}_{0:32} &\leftarrow \text{nprod}_{0:31} + (\text{RT}) \\ (\text{RT}) &\leftarrow \text{temp}_{1:32} \end{aligned}$$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is subtracted from the contents of RT and RT is updated with the low-order 32 bits of the result.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## nmacchws

Negative Multiply Accumulate Cross Halfword to Word Saturate Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>nmacchws</b>	RT, RA, RB	OE=0, Rc=0
<b>nmacchws.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmacchwso</b>	RT, RA, RB	OE=1, Rc=0
<b>nmacchwso.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	238	Rc
0	6	11	16	21 22		31

```
nprod0:31 ← -((RA)16:31 × (RB)0:15 signed)
temp0:32 ← nprod0:31 + (RT)
if ((nprod0 = RT0) ∧ (RT0 ≠ temp1)) then (RT) ← (RT0 || 31(¬RT0))
else (RT) ← temp1:32
```

The low-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is subtracted from the contents of RT.

If the result of the subtraction can be represented in 32 bits, then RT is updated with the low-order 32 bits of the result.

If the result of the subtraction cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the result is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the result is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>nmachhw</b>	RT, RA, RB	OE=0, Rc=0
<b>nmachhw.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmachhwo</b>	RT, RA, RB	OE=1, Rc=0
<b>nmachhwo.</b>	RT, RA, RB	OE=1, Rc=1



$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$  signed  
 $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$   
 $(RT) \leftarrow temp_{1:32}$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is subtracted from the contents of RT and RT is updated with the low-order 32 bits of the result.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## nmachhws

Negative Multiply Accumulate High Halfword to Word Saturate Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>nmachhws</b>	RT, RA, RB	OE=0, Rc=0
<b>nmachhws.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmachhwsO</b>	RT, RA, RB	OE=1, Rc=0
<b>nmachhwsO.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	110	Rc
0	6	11	16	21 22		31

$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$  signed  
 $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$   
if  $((nprod_0 = RT_0) \wedge (RT_0 \neq temp_1))$  then  $(RT) \leftarrow (RT_0 \parallel^{31}(\neg RT_0))$   
else  $(RT) \leftarrow temp_{1:32}$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The signed product is subtracted from the contents of RT.

If the result of the subtraction can be represented in 32 bits, then RT is updated with the low-order 32 bits of the result.

If the result of the subtraction cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the result is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the result is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.



**Preliminary**

<b>nmac1hw</b>	RT, RA, RB	OE=0, Rc=0
<b>nmac1hw.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmac1hwo</b>	RT, RA, RB	OE=1, Rc=0
<b>nmac1hwo.</b>	RT, RA, RB	OE=1, Rc=1



$$\text{nprod}_{0:31} \leftarrow -((\text{RA})_{16:31} \times (\text{RB})_{16:31}) \text{ signed}$$

$$\text{temp}_{0:32} \leftarrow \text{nprod}_{0:31} + (\text{RT})$$

$$(\text{RT}) \leftarrow \text{temp}_{1:32}$$

The low-order halfword of RA is multiplied by the low-order halfword of RB. The signed product is subtracted from the contents of RT and RT is updated with the low-order 32 bits of the result.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

**Architecture Note**

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

## nmaclhws

Negative Multiply Accumulate High Halfword to Word Saturate Signed  
PPC440x5 CPU Core User's Manual



Preliminary

<b>nmaclhws</b>	RT, RA, RB	OE=0, Rc=0
<b>nmaclhws.</b>	RT, RA, RB	OE=0, Rc=1
<b>nmaclhws0</b>	RT, RA, RB	OE=1, Rc=0
<b>nmaclhws0.</b>	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	494	Rc
0	6	11	16	21 22		31

```
nprod0:31 ← −((RA)16:31 × (RB)16:31) signed
temp0:32 ← nprod0:31 + (RT)
if ((nprod0 = RT0) ∧ (RT0 ≠ temp1)) then (RT) ← (RT0 ∥ 31(¬RT0))
else (RT) ← temp1:32
```

The low-order halfword of RA is multiplied by the low-order halfword of RB. The signed product is subtracted from the contents of RT.

If the result of the subtraction can be represented in 32 bits, then RT is updated with the low-order 32 bits of the result.

If the result of the subtraction cannot be represented in 32 bits, then RT is updated with a value which is “saturated” to the nearest representable value. That is, if the result is less than  $-2^{31}$ , then RT is updated with  $-2^{31}$ . Likewise, if the result is greater than  $2^{31} - 1$ , then RT is updated with  $2^{31} - 1$ .


### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

### Architecture Note

This instruction is implementation-specific and programs which use this instruction may not be portable to other PowerPC Book-E implementations. See *Instruction Set Portability* on page 244.

**nor** RA, RS, RB R<sub>c</sub>=0  
**nor.** RA, RS, RB R<sub>c</sub>=1



$$(RA) \leftarrow \neg((RS) \vee (RB))$$

The contents of register RS is ORed with the contents of register RB; the ones complement of the result is placed into register RA.


**Registers Altered**

- RA
- CR[CR0] if R<sub>c</sub> contains 1

Table 9-24. Extended Mnemonics for nor, nor.

Mnemonic	Operands	Function	Other Registers Altered
<b>not</b>	RA, RS	Complement register. $(RA) \leftarrow \neg(RS)$ <i>Extended mnemonic for  <b>nor RA,RS,RS</b></i>	
<b>not.</b>		<i>Extended mnemonic for  <b>nor. RA,RS,RS</b></i>	CR[CR0]

**or** RA, RS, RB Rc=0  
**or.** RA, RS, RB Rc=1



$$(RA) \leftarrow (RS) \vee (RB)$$

The contents of register RS is ORed with the contents of register RB; the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if Rc contains 1

Table 9-25. Extended Mnemonics for or, or.


Mnemonic	Operands	Function	Other Registers Altered
<b>mr</b>	RT, RS	Move register. (RT) ← (RS) <i>Extended mnemonic for or RT,RS,RS</i>	
<b>mr.</b>		<i>Extended mnemonic for or. RT,RS,RS</i>	CR[CR0]





Preliminary

**orc** RA, RS, RB R<sub>c</sub>=0  
**orc.** RA, RS, RB R<sub>c</sub>=1



$$(RA) \leftarrow (RS) \vee \neg(RB)$$

The contents of register RS is ORed with the ones complement of the contents of register RB; the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if R<sub>c</sub> contains 1


# ori

OR Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

**ori** RA, RS, IM



$$(RA) \leftarrow (RS) \vee (^{16}0 \parallel IM)$$

The IM field is extended to 32 bits by concatenating 16 0-bits on the left. Register RS is ORed with the extended IM field; the result is placed into register RA.

## Registers Altered


- RA

Table 9-26. Extended Mnemonics for ori

Mnemonic	Operands	Function	Other Registers Changed
<b>nop</b>		Preferred no-op; triggers optimizations based on no-ops. <i>Extended mnemonic for <b>ori 0,0,0</b></i>	



**oris** RA, RS, IM



$$(RA) \leftarrow (RS) \vee (IM \parallel 16_0)$$

The IM Field is extended to 32 bits by concatenating 16 0-bits on the right. Register RS is ORed with the extended IM field and the result is placed into register RA.

**Registers Altered**

- RA


## rfci

Return From Critical Interrupt  
PPC440x5 CPU Core User's Manual



Preliminary

## rfci



(PC) ← (CSRR0)  
(MSR) ← (CSRR1)

This instruction is used to return from a critical interrupt.

The program counter (PC) is restored with the contents of CSRR0 and the MSR is restored with the contents of CSRR1.

Instruction execution returns to the address contained in the PC.


### Registers Altered

- MSR

### Programming Note

Execution of this instruction is privileged and context-synchronizing (see *Context Synchronization* on page 80).

**rfi**



(PC) ← (SRR0)  
(MSR) ← (SRR1)

This instruction is used to return from a non-critical interrupt.

The program counter (PC) is restored with the contents of SRR0 and the MSR is restored with the contents of SRR1.

Instruction execution returns to the address contained in the PC.

**Registers Altered**

- MSR

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Execution of this instruction is privileged and context-synchronizing (see *Context Synchronization* on page 80).


## rfmci

Return From Machine Check Interrupt  
PPC440x5 CPU Core User's Manual



Preliminary

### rfmci



(PC) ← (MCSRR0)  
(MSR) ← (MCSRR1)

This instruction is used to return from a machine check interrupt.

The program counter (PC) is restored with the contents of MCSRR0 and the MSR is restored with the contents of MCSRR1.

Instruction execution returns to the address contained in the PC.

#### Registers Altered

- MSR

#### Programming Note

Execution of this instruction is privileged and context-synchronizing (see “Context Synchronization” on page 80).

**Preliminary**

**rlwimi** RA, RS, SH, MB, ME R<sub>c</sub>=0  
**rlwimi.** RA, RS, SH, MB, ME R<sub>c</sub>=1

20	RS	RA	SH	MB	ME	R <sub>c</sub>
0	6	11	16	21	26	31

$r \leftarrow \text{ROTL}((RS), SH)$   
 $m \leftarrow \text{MASK}(MB, ME)$   
 $(RA) \leftarrow (r \wedge m) \vee ((RA) \wedge \neg m)$

The contents of register RS are rotated left by the number of bit positions specified in the SH field. A mask is generated, having 1-bits starting at the bit position specified in the MB field and ending in the bit position specified by the ME field, with 0-bits elsewhere.

If the starting point of the mask is at a higher bit position than the ending point, the 1-bits portion of the mask wraps from the highest bit position back around to the lowest. The rotated data is inserted into register RA, in positions corresponding to the bit positions in the mask that contain a 1-bit.

**Registers Altered**

- RA
- CR[CR0] if R<sub>c</sub> contains 1


*Table 9-27. Extended Mnemonics for rlwimi, rlwimi.*

Mnemonic	Operands	Function	Other Registers Altered
<b>inslwi</b>	RA, RS, n, b	Insert from left immediate (n > 0). $(RA)_{b:b+n-1} \leftarrow (RS)_{0:n-1}$ <i>Extended mnemonic for</i> <b>rlwimi RA,RS,32-b,b,b+n-1</b>	
<b>inslwi.</b>		<i>Extended mnemonic for</i> <b>rlwimi. RA,RS,32-b,b,b+n-1</b>	CR[CR0]
<b>insrwi</b>	RA, RS, n, b	Insert from right immediate. (n > 0) $(RA)_{b:b+n-1} \leftarrow (RS)_{32-n:31}$ <i>Extended mnemonic for</i> <b>rlwimi RA,RS,32-b-n,b,b+n-1</b>	
<b>insrwi.</b>		<i>Extended mnemonic for</i> <b>rlwimi. RA,RS,32-b-n,b,b+n-1</b>	CR[CR0]

# rlwinm

Rotate Left Word Immediate then AND with Mask

**rlwinm** RA, RS, SH, MB, ME Rc=0  
**rlwinm.** RA, RS, SH, MB, ME Rc=1



$r \leftarrow \text{ROTL}((RS), SH)$   
 $m \leftarrow \text{MASK}(MB, ME)$   
 $(RA) \leftarrow r \wedge m$

The contents of register RS are rotated left by the number of bit positions specified in the SH field. A mask is generated, having 1-bits starting at the bit position specified in the MB field and ending in the bit position specified by the ME field with 0-bits elsewhere.

If the starting point of the mask is at a higher bit position than the ending point, the 1-bits portion of the mask wraps from the highest bit position back around to the lowest. The rotated data is ANDed with the generated mask; the result is placed into register RA.

### Registers Altered

- RA
- CR[CR0] if Rc contains 1

Table 9-28. Extended Mnemonics for rlwinm, rlwinm.

Mnemonic	Operands	Function	Other Registers Altered
<b>clrlwi</b>	RA, RS, n	Clear left immediate. ( $n < 32$ ) $(RA)_{0:n-1} \leftarrow n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,0,n,31</b>	
<b>clrlwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,0,n,31</b>	CR[CR0]
<b>clrlslwi</b>	RA, RS, b, n	Clear left and shift left immediate. $(n \leq b < 32)$ $(RA)_{b-n:31-n} \leftarrow (RS)_{b:31}$ $(RA)_{32-n:31} \leftarrow n_0$ $(RA)_{0:b-n-1} \leftarrow b-n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,b-n,31-n</b>	
<b>clrlslwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,b-n,31-n</b>	CR[CR0]
<b>clrrwi</b>	RA, RS, n	Clear right immediate. ( $n < 32$ ) $(RA)_{32-n:31} \leftarrow n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,0,0,31-n</b>	
<b>clrrwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,0,0,31-n</b>	CR[CR0]



Table 9-28. Extended Mnemonics for rlwinm, rlwinm. (continued)

Mnemonic	Operands	Function	Other Registers Altered
<b>extlwi</b>	RA, RS, n, b	Extract and left justify immediate. ( $n > 0$ ) $(RA)_{0:n-1} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{n:31} \leftarrow 32-n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,b,0,n-1</b>	
<b>extlwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,b,0,n-1</b>	CR[CR0]
<b>extrwi</b>	RA, RS, n, b	Extract and right justify immediate. ( $n > 0$ ) $(RA)_{32-n:31} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{0:31-n} \leftarrow 32-n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,b+n,32-n,31</b>	
<b>extrwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,b+n,32-n,31</b>	CR[CR0]
<b>rotlwi</b>	RA, RS, n	Rotate left immediate. $(RA) \leftarrow \text{ROTL}((RS), n)$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,0,31</b>	
<b>rotlwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,0,31</b>	CR[CR0]
<b>rotrwi</b>	RA, RS, n	Rotate right immediate. $(RA) \leftarrow \text{ROTR}((RS), 32-n)$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,32-n,0,31</b>	
<b>rotrwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,32-n,0,31</b>	CR[CR0]
<b>slwi</b>	RA, RS, n	Shift left immediate. ( $n < 32$ ) $(RA)_{0:31-n} \leftarrow (RS)_{n:31}$ $(RA)_{32-n:31} \leftarrow n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,0,31-n</b>	
<b>slwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,0,31-n</b>	CR[CR0]
<b>srwi</b>	RA, RS, n	Shift right immediate. ( $n < 32$ ) $(RA)_{n:31} \leftarrow (RS)_{0:31-n}$ $(RA)_{0:n-1} \leftarrow n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,32-n,n,31</b>	
<b>srwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,32-n,n,31</b>	CR[CR0]

## rlwnm

Rotate Left Word then AND with Mask  
PPC440x5 CPU Core User's Manual



Preliminary

<b>rlwnm</b>	RA, RS, RB, MB, ME	Rc=0
<b>rlwnm.</b>	RA, RS, RB, MB, ME	Rc=1

23	RS	RA	RB	MB	ME	Rc
0	6	11	16	21	26	31

$$r \leftarrow \text{ROTL}((RS), (RB)_{27:31})$$
$$m \leftarrow \text{MASK}(MB, ME)$$
$$(RA) \leftarrow r \wedge m$$

The contents of register RS are rotated left by the number of bit positions specified by the contents of register RB<sub>27:31</sub>. A mask is generated, having 1-bits starting at the bit position specified in the MB field and ending in the bit position specified by the ME field with 0-bits elsewhere.


If the starting point of the mask is at a higher bit position than the ending point, the ones portion of the mask wraps from the highest bit position back to the lowest. The rotated data is ANDed with the generated mask and the result is placed into register RA.

### Registers Altered

- RA
- CR[CR0] if Rc contains 1

Table 9-29. Extended Mnemonics for rlwnm, rlwnm.

Mnemonic	Operands	Function	Other Registers Altered
<b>rotlw</b>	RA, RS, RB	Rotate left. $(RA) \leftarrow \text{ROTL}((RS), (RB)_{27:31})$ <i>Extended mnemonic for</i> <b>rlwnm RA,RS,RB,0,31</b>	
<b>rotlw.</b>		<i>Extended mnemonic for</i> <b>rlwnm. RA,RS,RB,0,31</b>	CR[CR0]

**sc**


$SRR1 \leftarrow MSR$   
 $SRR0 \leftarrow 4 + \text{address of } \mathbf{sc} \text{ instruction}$   
 $PC \leftarrow IVPR_{0:15} \parallel IVOR8_{16:27} \parallel 40$   
 $MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] \leftarrow 0$

A System Call exception is generated, and a System Call interrupt occurs (see *System Call Interrupt* on page 184 for more information on System Call interrupts). The contents of the MSR are copied into SRR1 and (4 + address of **sc** instruction) is placed into SRR0.

The program counter (PC) is then loaded with the interrupt vector address. The interrupt vector address is formed by concatenating the high halfword of the Interrupt Vector Prefix Register (IVPR), bits 16:27 of the Interrupt Vector Offset Register 8 (IVOR8), and 0b0000.

The MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] bits are set to 0.

Program execution continues at the new address in the PC.

**Registers Altered**

- SRR0
- SRR1
- MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS]

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Execution of this instruction is context-synchronizing (see *Context Synchronization* on page 80).

## slw


Shift Left Word

PPC440x5 CPU Core User's Manual



Preliminary

**slw** RA, RS, RB Rc=0  
**slw.** RA, RS, RB Rc=1



```
n ← (RB)26:31
r ← ROTL((RS), n)
if n < 32 then
    m ← MASK(0, 31 – n)
else
    m ← 320
(RA) ← r ∧ m
```

The contents of register RS are shifted left by the number of bits specified by the contents of register RB<sub>26:31</sub>. Bits shifted left out of the most significant bit are lost, and 0-bits fill vacated bit positions on the right. The result is placed into register RA.


Note that if RB<sub>26</sub> = 1, then the shift amount is 32 bits or more, and thus all bits are shifted out such that register RA is set to zero.

### Registers Altered

- RA
- CR[CR0] if Rc contains 1



<b>sraw</b>	RA, RS, RB	Rc=0
<b>sraw.</b>	RA, RS, RB	Rc=1



```

n ← (RB)26:31
r ← ROTL((RS), 32 - n)
if n < 32 then
  m ← MASK(n, 31)
else
  m ← 320
s ← (RS)0
(RA) ← (r ∧ m) ∨ (32s ∧ ¬m)
XER[CA] ← s ∧ ((r ∧ ¬m) ≠ 0)

```

The contents of register RS are shifted right by the number of bits specified the contents of register RB<sub>26:31</sub>. Bits shifted out of the least significant bit are lost. Bit 0 of Register RS is replicated to fill the vacated positions on the left. The result is placed into register RA.

If register RS contains a negative number and any 1-bits were shifted out of the least significant bit position, XER[CA] is set to 1; otherwise, it is set to 0.

Note that if RB<sub>26</sub> = 1, then the shift amount is 32 bits or more, and thus all bits are shifted out such that register RA and XER[CA] are set to bit 0 of register RS.

**Registers Altered**

- RA
- XER[CA]
- CR[CR0] if Rc contains 1


## srawi

Shift Right Algebraic Word Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

<b>srawi</b>	RA, RS, SH	Rc=0
<b>srawi.</b>	RA, RS, SH	Rc=1



```
n ← SH
r ← ROTL((RS), 32 - n)
m ← MASK(n, 31)
s ← (RS)0
(RA) ← (r ∧ m) ∨ (32s ∧ ¬m)
XER[CA] ← s ∧ ((r ∧ ¬m)≠0)
```

The contents of register RS are shifted right by the number of bits specified in the SH field. Bits shifted out of the least significant bit are lost. Bit RS<sub>0</sub> is replicated to fill the vacated positions on the left. The result is placed into register RA.


If register RS contains a negative number and any 1-bits were shifted out of the least significant bit position, XER[CA] is set to 1; otherwise, it is set to 0.

### Registers Altered

- RA
- XER[CA]
- CR[CR0] if Rc contains 1



**srw** RA, RS, RB R<sub>c</sub>=0  
**srw.** RA, RS, RB R<sub>c</sub>=1



```

n ← (RB)26:31
r ← ROTL((RS), 32 - n)
if n < 32 then
  m ← MASK(n, 31)
else
  m ← 320
(RA) ← r ∧ m

```

The contents of register RS are shifted right by the number of bits specified the contents of register RB<sub>26:31</sub>. Bits shifted right out of the least significant bit are lost, and 0-bits fill the vacated bit positions on the left. The result is placed into register RA.

Note that if RB<sub>26</sub> = 1, then the shift amount is 32 bits or more, and thus all bits are shifted out such that register RA is set to zero.

#### Registers Altered

- RA
- CR[CR0] if R<sub>c</sub> contains 1

## stb


Store Byte

PPC440x5 CPU Core User's Manual



Preliminary

**stb**            RS, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$

$MS(EA, 1) \leftarrow (RS)_{24:31}$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant byte of register RS is stored into the byte at the EA.


### Registers Altered

- None





**stbu**            RS, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $MS(EA, 1) \leftarrow (RS)_{24:31}$   
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant byte of register RS is stored into the byte at the EA.

The EA is placed into register RA.

**Registers Altered**

- RA

**Invalid Instruction Forms**

RA = 0


## stbx

Store Byte with Update Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**stbx**            RS, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $MS(EA, 1) \leftarrow (RS)_{24:31}$   
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant byte of register RS is stored into the byte at the EA.

The EA is placed into register RA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered


- RA

### Invalid Instruction Forms

- Reserved fields

RA = 0

**stbx**            RS, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$MS(EA, 1) \leftarrow (RS)_{24:31}$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant byte of register RS is stored into the byte at the EA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields


## sth

Store Halfword  
PPC440x5 CPU Core User's Manual



Preliminary

**sth**            RS, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $MS(EA, 2) \leftarrow (RS)_{16:31}$


An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0 and is the contents of register RA otherwise.

The least significant halfword of register RS is stored into the halfword at the EA in main storage.

### Registers Altered

- None

**sthbrx**      RS, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$MS(EA, 2) \leftarrow \text{BYTE\_REVERSE}((RS)_{16:31})$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0 and is the contents of register RA otherwise.

The least significant halfword of register RS is byte-reversed from the default byte ordering for the memory page referenced by the EA. The resulting halfword is stored at the EA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Byte ordering is generally controlled by the Endian (E) storage attribute (see *Memory Management* on page 129). The store byte reverse instructions provide a mechanism for data to be stored to a memory page using the opposite byte ordering from that specified by the Endian storage attribute.


## sth

Store Halfword with Update  
PPC440x5 CPU Core User's Manual



Preliminary

**sth**            RS, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$   
 $MS(EA, 2) \leftarrow (RS)_{16:31}$   
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant halfword of register RS is stored into the halfword at the EA.

The EA is placed into register RA.


### Registers Altered

- RA

### Invalid Instruction Forms

RA = 0

**sthux**            RS, RA, RB



$$\begin{aligned}
 EA &\leftarrow (RA|0) + (RB) \\
 MS(EA, 2) &\leftarrow (RS)_{16:31} \\
 (RA) &\leftarrow EA
 \end{aligned}$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant halfword of register RS is stored into the halfword at the EA.

The EA is placed into register RA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RA

**Invalid Instruction Forms**

- Reserved fields
- RA = 0


## sthx

Store Halfword Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**sthx**            RS, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $MS(EA, 2) \leftarrow (RS)_{16:31}$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The least significant halfword of register RS is stored into the halfword at the EA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered


- None

### Invalid Instruction Forms

- Reserved fields



**stmw**            RS, D(RA)



```
EA ← (RA|0) + EXTS(D)
r ← RS
do while r ≤ 31
  MS(EA, 4) ← (GPR(r))
  r ← r + 1
  EA ← EA + 4
```

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The contents of a series of consecutive registers, starting with register RS and continuing through GPR(31), are stored into consecutive words starting at the EA.

**Registers Altered**

- None

**Programming Note**

This instruction can be restarted, meaning that it could be interrupted after having already stored some of the register values to memory, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been stored prior to the interrupt will be stored a second time.

## stswi


Store String Word Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

Store String Word Immediate

**stswi**            RS, RA, NB



```
EA ← (RA|0)
if NB = 0 then
    n ← 32
else
    n ← NB
r ← RS - 1
i ← 0
do while n > 0
    if i = 0 then
        r ← r + 1
    if r = 32 then
        r ← 0
    MS(EA,1) ← (GPR(r)i:i+7)
    i ← i + 8
    if i = 32 then
        i ← 0
    EA ← EA + 1
    n ← n - 1
```

An effective address (EA) is determined by the RA field. If the RA field contains 0, the EA is 0; otherwise, the EA is the contents of register RA.

A byte count is determined by the NB field. If the NB field contains 0, the byte count is 32; otherwise, the byte count is the contents of the NB field.

The contents of a series of consecutive GPRs (starting with register RS, continuing through GPR(31) and wrapping to GPR(0) as necessary, and continuing to the final byte count) are stored, starting at the EA. The bytes in each GPR are accessed starting with the most significant byte. The byte count determines the number of transferred bytes.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- None


### Invalid Instruction Forms

- Reserved fields

### Programming Note

This instruction can be restarted, meaning that it could be interrupted after having already stored some of the register values to memory, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been stored prior to the interrupt will be stored a second time.

**stswx**      RS, RA, RB



```

EA ← (RA|0) + (RB)
n ← XER[TBC]
r ← RS - 1
i ← 0
do while n > 0
  if i = 0 then
    r ← r + 1
  if r = 32 then
    r ← 0
  MS(EA, 1) ← (GPR(r)i:i+7)
  i ← i + 8
  if i = 32 then
    i ← 0
  EA ← EA + 1
  n ← n - 1

```

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

A byte count is contained in XER[TBC].

The contents of a series of consecutive GPRs (starting with register RS, continuing through GPR(31) and wrapping to GPR(0) as necessary, and continuing to the final byte count) are stored, starting at the EA. The bytes in each GPR are accessed starting with the most significant byte. The byte count determines the number of transferred bytes.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

This instruction can be restarted, meaning that it could be interrupted after having already stored some of the register values to memory, and then re-executed from the beginning (after returning from the interrupt), in which case the registers which had already been stored prior to the interrupt will be stored a second time.

If XER[TBC] = 0, no GPRs are stored to memory, and **stswx** is treated as a no-op. Furthermore, if the EA is such that a Data Storage, Data TLB Error, or Data Address Compare Debug exception occurs, **stswx** is treated as a no-op and no interrupt occurs as a result of the exception.

## stw


Store Word

PPC440x5 CPU Core User's Manual



Preliminary

**stw**            RS, D(RA)



$EA \leftarrow (RA|0) + \text{EXTS}(D)$

$MS(EA, 4) \leftarrow (RS)$


An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The contents of register RS are stored at the EA.

### Registers Altered

- None

**stwbrx**      RS, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$MS(EA, 4) \leftarrow \text{BYTE\_REVERSE}((RS)_{0:31})$$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0 and is the contents of register RA otherwise.

The word in register RS is byte-reversed from the default byte ordering for the memory page referenced by the EA. The resulting word is stored at the EA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- None

**Invalid Instruction Forms**

- Reserved fields

**Programming Note**

Byte ordering is generally controlled by the Endian (E) storage attribute (see *Memory Management* on page 129). The store byte reverse instructions provide a mechanism for data to be stored to a memory page using the opposite byte ordering from that specified by the Endian storage attribute.


## stwcx.

Store Word Conditional Indexed  
PPC440x5 CPU Core User's Manual



Preliminary

**stwcx.** RS, RA, RB



```
EA ← (RA|0) + (RB)
if RESERVE = 1 then
    MS(EA, 4) ← (RS)
    RESERVE ← 0
    (CR[CR0]) ← 20 || 1 || XER[SO]
else
    (CR[CR0]) ← 20 || 0 || XER[SO]
```

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

If the reservation bit contains 1 when the instruction is executed, the contents of register RS are stored into the word at the EA and the reservation bit is cleared. If the reservation bit contains 0 when the instruction is executed, no store operation is performed.

CR[CR0] is set as follows:

- CR[CR0]<sub>0:1</sub> are cleared
- CR[CR0]<sub>2</sub> is set to indicate whether or not the store was performed (1 indicates that it was)
- CR[CR0]<sub>3</sub> is set to the contents of the XER[SO] bit

### Registers Altered

- CR[CR0]

### Programming Notes

The **lwarx** and **stwcx.** instructions are typically paired in a loop, as shown in the following example, to create the effect of an atomic operation to a memory area used as a semaphore between multiple processes. Only **lwarx** can set the reservation bit to 1. **stwcx.** sets the reservation bit to 0 upon its completion, whether or not **stwcx.** actually stored (RS) to memory. CR[CR0]<sub>2</sub> must be examined to determine whether (RS) was sent to memory.

```
loop: lwarx      # read the semaphore from memory; set reservation
      "alter"   # change the semaphore bits in the register as required
      stwcx.    # attempt to store the semaphore; reset reservation
      bne loop  # some other process intervened and cleared the reservation prior to the above
              # stwcx.; try again
```

The PowerPC Book-E architecture specifies that the EA for the **lwarx** instruction must be word-aligned (that is, a multiple of 4 bytes); otherwise, the result is undefined. Although the PPC440x5 will execute **stwcx.** regardless of the EA alignment, in order for the operation of the pairing of **lwarx** and **stwcx.** to produce the desired result, software must ensure that the EA for both instructions is word-aligned. This requirement is due to the manner in which misaligned storage accesses may be broken up into separate, aligned accesses by the PPC440x5.



The PowerPC Book-E architecture also specifies that it is implementation-dependent as to whether a Data Storage, Data TLB Error, Alignment, or Debug interrupt occurs when the reservation bit is off at the time of execution of an **stwcx.** instruction, and when the conditions are such that a non-**stwcx.** store-type storage access instruction would have resulted in such an interrupt. The PPC440x5 implements **stwcx.** such that Data Storage and Debug (DAC and/or DVC exception type) interrupts do not occur when the reservation bit is off at the time of execution of the **stwcx.** Instead, the **stwcx.** instruction completes without causing the interrupt and without storing to memory, and CR[CR0] is updated to indicate the failure of the **stwcx.**

On the other hand, the PPC440x5 causes a Data TLB Error interrupt if a Data TLB Miss exception occurs during due to the execution of a **stwcx.** instruction, regardless of the state of the reservation. Similarly, the PPC440x5 causes an Alignment interrupt if the EA of the **stwcx.** operand is not word-aligned when CCR0[FLSTA] is 1, regardless of the state of the reservation (see *Core Configuration Register 0 (CCR0)* on page 105 for more information on the Force Load/Store Alignment function).

## stwu


Store Word with Update

PPC440x5 CPU Core User's Manual



Preliminary

**stwu**            RS, D(RA)



$EA \leftarrow (RA|0) + \text{EXTS}(D)$

$MS(EA, 4) \leftarrow (RS)$

$(RA) \leftarrow EA$

An effective address (EA) is formed by adding a displacement to a base address. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The contents of register RS are stored into the word at the EA.

The EA is placed into register RA.

### Registers Altered


- RA

### Invalid Instruction Forms

RA = 0



**stwux**      RS, RA, RB



$EA \leftarrow (RA|0) + (RB)$   
 $MS(EA, 4) \leftarrow (RS)$   
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The contents of register RS are stored into the word at the EA.

The EA is placed into register RA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- RA

**Invalid Instruction Forms**

- Reserved fields
- RA = 0

## stwx


Store Word Indexed

PPC440x5 CPU Core User's Manual



Preliminary

**stwx**            RS, RA, RB



$EA \leftarrow (RA|0) + (RB)$

$MS(EA,4) \leftarrow (RS)$

An effective address (EA) is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 when the RA field is 0, and is the contents of register RA otherwise.

The contents of register RS are stored into the word at the EA.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.


### Registers Altered

- None

### Invalid Instruction Forms

- Reserved fields

<b>subf</b>	RT, RA, RB	OE=0, Rc=0
<b>subf.</b>	RT, RA, RB	OE=0, Rc=1
<b>subfo</b>	RT, RA, RB	OE=1, Rc=0
<b>subfo.</b>	RT, RA, RB	OE=1, Rc=1



$$(RT) \leftarrow \neg(RA) + (RB) + 1$$

The sum of the ones complement of register RA, register RB, and 1 is stored into register RT.

### Registers Altered

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

Table 9-30. Extended Mnemonics for *subf*, *subf.*, *subfo*, *subfo.*

Mnemonic	Operands	Function	Other Registers Altered
<b>sub</b>	RT, RA, RB	Subtract (RB) from (RA). $(RT) \leftarrow \neg(RB) + (RA) + 1.$ <i>Extended mnemonic for</i> <b>subf RT,RB,RA</b>	
<b>sub.</b>		<i>Extended mnemonic for</i> <b>subf. RT,RB,RA</b>	CR[CR0]
<b>subfo</b>		<i>Extended mnemonic for</i> <b>subfo RT,RB,RA</b>	XER[SO, OV]
<b>subfo.</b>		<i>Extended mnemonic for</i> <b>subfo. RT,RB,RA</b>	CR[CR0] XER[SO, OV]

## subfc

Subtract From Carrying

PPC440x5 CPU Core User's Manual



Preliminary

<b>subfc</b>	RT, RA, RB	OE=0, Rc=0
<b>subfc.</b>	RT, RA, RB	OE=0, Rc=1
<b>subfco</b>	RT, RA, RB	OE=1, Rc=0
<b>subfco.</b>	RT, RA, RB	OE=1, Rc=1

31	RT	RA	RB	OE	8	Rc
0	6	11	16	21 22		31

```
(RT) ← ¬(RA) + (RB) + 1
if ¬(RA) + (RB) + 1 ≥ 232 - 1 then
  XER[CA] ← 1
else
  XER[CA] ← 0
```

The sum of the ones complement of register RA, register RB, and 1 is stored into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the subtract operation.

### Registers Altered


- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

Table 9-31. Extended Mnemonics for *subfc*, *subfc.*, *subfco*, *subfco.*

Mnemonic	Operands	Function	Other Registers Altered
<b>subc</b>	RT, RA, RB	Subtract (RB) from (RA). $(RT) \leftarrow \neg(RB) + (RA) + 1$ . Place carry-out in XER[CA]. <i>Extended mnemonic for subfc RT,RB,RA</i>	
<b>subc.</b>		<i>Extended mnemonic for subfco RT,RB,RA</i>	CR[CR0]
<b>subfco</b>		<i>Extended mnemonic for subfc RT,RB,RA</i>	XER[SO, OV]
<b>subfco.</b>		<i>Extended mnemonic for subfco RT,RB,RA</i>	CR[CR0] XER[SO, OV]



<b>subfe</b>	RT, RA, RB	OE=0, Rc=0
<b>subfe.</b>	RT, RA, RB	OE=0, Rc=1
<b>subfeo</b>	RT, RA, RB	OE=1, Rc=0
<b>subfeo.</b>	RT, RA, RB	OE=1, Rc=1



```

(RT) ← ¬(RA) + (RB) + XER[CA]
if ¬(RA) + (RB) + XER[CA] > 232 - 1 then
  XER[CA] ← 1
else
  XER[CA] ← 0

```

The sum of the ones complement of register RA, register RB, and XER[CA] is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the subtract operation.

**Registers Altered**

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1


## subfic

Subtract From Immediate Carrying  
PPC440x5 CPU Core User's Manual



Preliminary

**subfic**      RT, RA, IM



```
(RT) ← ¬(RA) + EXTS(IM) + 1
if ¬(RA) + EXTS(IM) + 1 > 232 - 1 then
    XER[CA] ← 1
else
    XER[CA] ← 0
```

The sum of the ones complement of RA, the IM field sign-extended to 32 bits, and 1 is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the subtract operation.


### Registers Altered

- RT
- XER[CA]



**Preliminary**

<b>subfme</b>	RT, RA	OE=0, Rc=0
<b>subfme.</b>	RT, RA	OE=0, Rc=1
<b>subfmeo</b>	RT, RA	OE=1, Rc=0
<b>subfmeo.</b>	RT, RA	OE=1, Rc=1



```

(RT) ← ¬(RA) - 1 + XER[CA]
if ¬(RA) + 0xFFFF FFFF + XER[CA] > 232 - 1 then
    XER[CA] ← 1
else
    XER[CA] ← 0

```

The sum of the ones complement of register RA, -1, and XER[CA] is placed into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the subtract operation.

**Registers Altered**

- RT
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1
- XER[CA]

**Invalid Instruction Forms**

- Reserved fields


## subfze

Subtract from Zero Extended  
PPC440x5 CPU Core User's Manual



Preliminary

<b>subfze</b>	RT, RA	OE=0, Rc=0
<b>subfze.</b>	RT, RA	OE=0, Rc=1
<b>subfzeo</b>	RT, RA	OE=1, Rc=0
<b>subfzeo.</b>	RT, RA	OE=1, Rc=1



```
(RT) ← ¬(RA) + XER[CA]
if ¬(RA) + XER[CA] > 232 - 1 then
    XER[CA] ← 1
else
    XER[CA] ← 0
```

The sum of the ones complement of register RA and XER[CA] is stored into register RT.

XER[CA] is set to a value determined by the unsigned magnitude of the result of the subtract operation.


### Registers Altered

- RT
- XER[CA]
- CR[CR0] if Rc contains 1
- XER[SO, OV] if OE contains 1

### Invalid Instruction Forms

- Reserved fields



**tlbre**      RT, RA, WS


```

tlbentry ← TLB[(RA)26:31]
if WS = 0
  (RT)0:27 ← tlbentry[EPN,V,TS,SIZE]
  if CCR0[CRPE] = 0
    (RT)28:31 ← 40
  else
    (RT)28:31 ← TPAR
  MMUCR[STID] ← tlbentry[TID]
else if WS = 1
  (RT)0:21 ← tlbentry[RPN]
  if CCR0[CRPE] = 0
    (RT)22:23 ← 20
  else
    (RT)22:23 ← PAR1
  (RT)24:27 ← 40
  (RT)28:31 ← tlbentry[ERPNI]
else if WS = 2
  if CCR0[CRPE] = 0
    (RT)0:1 ← 20
  else
    (RT)0:1 ← PAR2
  (RT)2:15 ← 140
  (RT)16:24 ← tlbentry[U0,U1,U2,U3,W,I,M,G,E]
  (RT)25 ← 0
  (RT)26:31 ← tlbentry[UX,UW,UR,SX,SW,SR]
else (RT), MMUCR[STID] ← undefined

```

The contents of the specified portion of the selected TLB entry are placed into register RT (and also MMUCR[STID] if WS = 0).

The parity bits in the TLB entry (TPAR, PAR1, and PAR2) are placed into the register RT if and only if the Cache Read Parity Enable bit, CCR0[CRPE], is set to 1.

The contents of RA are used as an index into the TLB. If this value is greater than the index of the highest numbered TLB entry (63), the results are undefined.

The WS field specifies which portion of the TLB entry is placed into RT. If WS = 0, the TID field of the selected TLB entry is read into MMUCR[STID]. See *Memory Management* on page 129 for descriptions of the TLB entry fields.

If the value of the WS field is greater than 2, the instruction form is invalid and the result is undefined.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

## **tlbre**

TLB Read Entry

PPC440x5 CPU Core User's Manual



Preliminary

### **Registers Altered**

- RT
- MMUCR[STID] (if WS = 0)

### **Invalid Instruction Forms**


- Reserved fields
- Invalid WS value

### **Programming Notes**

Execution of this instruction is privileged.

The PPC440x5 core does not automatically synchronize the context of the MMUCR[STID] field between a **tlbre** instruction which updates the field, and a **tlbsx**[.] instruction which uses it as a source operand. Therefore, software must execute an **isync** instruction between the execution of a **tlbre** instruction and a subsequent **tlbsx**[.] instruction to ensure that the **tlbsx**[.] instruction will use the new value of MMUCR[STID]. On the other hand, the PPC440x5 core *does* automatically synchronize the context of MMUCR[STID] between **tlbre** and **tlbwe**, as well as between **tlbre** and **mfspr** which specifies the MMUCR as the source SPR, so no **isync** is required in these cases.

<b>tlbsx</b>	RT, RA, RB	Rc=0
<b>tlbsx.</b>	RT, RA, RB	Rc=1



```

EA ← (RA|0) + (RB)
if Rc = 1
  CR[CR0]0 ← 0
  CR[CR0]1 ← 0
  CR[CR0]3 ← XER[SO]
if Valid TLB entry matching EA and MMUCR[STID,STS] is in the TLB then
  (RT) ← Index of matching TLB Entry
  if Rc = 1
    CR[CR0]2 ← 1
else
  (RT) ← Undefined
  if Rc = 1
    CR[CR0]2 ← 0
  
```

An effective address is formed by adding an index to a base address. The index is the contents of register RB. The base address is 0 if the RA field is 0 and is the contents of register RA otherwise.

The TLB is searched for a valid entry which translates EA and MMUCR[STID,STS]. See *Memory Management* on page 129 for descriptions of the TLB fields and how they participate in the determination of a match. If a matching entry is found, its index (0 - 63) is placed into bits 26:31 of RT, and bits 0:25 are set to 0. If no match is found, the contents of RT are undefined.

The record bit (Rc) specifies whether the results of the search will affect CR[CR0] as shown above, such that CR[CR0]<sub>2</sub> can be tested if there is a possibility that the search may fail.

### Registers Altered

- CR[CR0] if Rc contains 1

### Invalid Instruction Forms

- None

### Programming Notes

Execution of this instruction is privileged.

The PPC440x5 core does not automatically synchronize the context of the MMUCR[STID] field between a **tlbre** instruction which updates the field, and a **tlbsx**[.] instruction which uses it as a source operand. Therefore, software must execute an **isync** instruction between the execution of a **tlbre** instruction and a subsequent **tlbsx**[.] instruction to ensure that the **tlbsx**[.] instruction will use the new value of MMUCR[STID]. On the other hand, the PPC440x5 core *does* automatically synchronize the context of MMUCR[STID] between **tlbre** and **tlbwe**, as well as between **tlbre** and **mfspr** which specifies the MMUCR as the source SPR, so no **isync** is required in these cases.

## tlbsync


TLB Synchronize

PPC440x5 CPU Core User's Manual



Preliminary

### tlbsync



The **tlbsync** instruction is provided by the PowerPC Book-E architecture to support synchronization of TLB operations between processors in a coherent multi-processor system. Since the PPC440x5 core does not support coherent multi-processing, this instruction performs no operation, and is provided only to facilitate code portability.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

#### Registers Altered

- None


#### Invalid Instruction Forms

- Reserved fields

#### Programming Note

This instruction is privileged. Translation is not required to be active during the execution of this instruction.

Since the PPC440x5 core does not support tightly-coupled multiprocessor systems, **tlbsync** performs no operation.

**tlbwe**      RS, RA, WS


```

tlbentry ← TLB[(RA)26:31]
if WS = 0
  tlbentry[EPN,V,TS,SIZE] ← (RS)0:27
  tlbentry[TID] ← MMUCR[STID]
else if WS = 1
  tlbentry[RPN] ← (RS)0:21
  tlbentry[ERPNI] ← (RS)28:31
else if WS = 2
  tlbentry[U0,U1,U2,U3,W,I,M,G,E] ← (RS)16:24
  tlbentry[UX,UW,UR,SX,SW,SR] ← (RS)26:31
else tlbentry ← undefined
  
```

The contents of the specified portion of the selected TLB entry are replaced with the contents of register RS (and also MMUCR[STID] if WS = 0).

Parity check bits are automatically calculated and stored in the TLB entry as the tlbwe is executed. The contents of the RS register in the TPAR, PAR1, and PAR2 fields (for WS=0,1,or 2, respectively) is ignored by tlbwe; the parity is calculated from the other data bits being written to the TLB entry.

The contents of RA are used as an index into the TLB. If this value is greater than the index of the highest numbered TLB entry (63), the results are undefined.

The WS field specifies which portion of the TLB entry is replaced by the contents of RS. If WS = 0, the TID field of the selected TLB entry is replaced by the value in MMUCR[STID]. See *Memory Management* on page 129 for descriptions of the TLB entry fields.

If the value of the WS field is greater than 2, the instruction form is invalid and the result is undefined.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

#### Registers Altered

- None


#### Invalid Instruction Forms

- Reserved fields
- Invalid WS value

#### Programming Note

Execution of this instruction is privileged.

tw TO, RA, RB



```

if ( ((RA) < (RB) ^ TO0 = 1) ∨
      ((RA) > (RB) ^ TO1 = 1) ∨
      ((RA) = (RB) ^ TO2 = 1) ∨
      ((RA) u (RB) ^ TO3 = 1) ∨
      ((RA) u (RB) ^ TO4 = 1) )
  SRR0 ← address of tw instruction
  SRR1 ← MSR
  ESR[PTR] ← 1 (other bits cleared)
  MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] ← 90
  PC ← IVPR0:15 || IVOR616:27 || 40
else no operation

```

Register RA is compared with register RB. If any comparison condition enabled by the TO field is true, a Trap exception type Program interrupt occurs as follows (see *Program Interrupt* on page 180 for more information on Program interrupts). The contents of the MSR are copied into SRR1 and the address of the tw instruction) is placed into SRR0. ESR[PTR] is set to 1 and the other bits ESR bits cleared to indicate the type of exception causing the Program interrupt.

The program counter (PC) is then loaded with the interrupt vector address. The interrupt vector address is formed by concatenating the high halfword of the Interrupt Vector Prefix Register (IVPR), bits 16:27 of the Interrupt Vector Offset Register 6 (IVOR6), and 0b0000.

MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] are set to 0.

Program execution continues at the new address in the PC.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- SRR0 (if trap condition is met)
- SRR1 (if trap condition is met)
- MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] (if trap condition is met)
- ESR (if trap condition is met)

### Invalid Instruction Forms

- Reserved fields

### Programming Notes

This instruction can be inserted into the execution stream by a debugger to implement breakpoints, and is not typically used by application code.

The enabling of trap debug events may affect the interrupt type caused by the execution of **tw** instruction. Specifically, trap instructions may be enabled to cause Debug interrupts instead of Program interrupts. See *Trap (TRAP) Debug Event* on page 228 for more details.

 Table 9-32. Extended Mnemonics for **tw**

Mnemonic	Operands	Function	Other Registers Altered
<b>trap</b>		Trap unconditionally. <i>Extended mnemonic for</i> <b>tw 31,0,0</b>	
<b>tweq</b>	RA, RB	Trap if (RA) equal to (RB). <i>Extended mnemonic for</i> <b>tw 4,RA,RB</b>	
<b>twge</b>	RA, RB	Trap if (RA) greater than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 12,RA,RB</b>	
<b>twgt</b>	RA, RB	Trap if (RA) greater than (RB). <i>Extended mnemonic for</i> <b>tw 8,RA,RB</b>	
<b>twle</b>	RA, RB	Trap if (RA) less than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 20,RA,RB</b>	
<b>twlge</b>	RA, RB	Trap if (RA) logically greater than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 5,RA,RB</b>	
<b>twlgt</b>	RA, RB	Trap if (RA) logically greater than (RB). <i>Extended mnemonic for</i> <b>tw 1,RA,RB</b>	
<b>twlle</b>	RA, RB	Trap if (RA) logically less than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 6,RA,RB</b>	
<b>twllt</b>	RA, RB	Trap if (RA) logically less than (RB). <i>Extended mnemonic for</i> <b>tw 2,RA,RB</b>	
<b>twlng</b>	RA, RB	Trap if (RA) logically not greater than (RB). <i>Extended mnemonic for</i> <b>tw 6,RA,RB</b>	
<b>twlnl</b>	RA, RB	Trap if (RA) logically not less than (RB). <i>Extended mnemonic for</i> <b>tw 5,RA,RB</b>	
<b>twlt</b>	RA, RB	Trap if (RA) less than (RB). <i>Extended mnemonic for</i> <b>tw 16,RA,RB</b>	
<b>twne</b>	RA, RB	Trap if (RA) not equal to (RB). <i>Extended mnemonic for</i> <b>tw 24,RA,RB</b>	
<b>twng</b>	RA, RB	Trap if (RA) not greater than (RB). <i>Extended mnemonic for</i> <b>tw 20,RA,RB</b>	


**tw**

Trap Word

*Table 9-32. Extended Mnemonics for tw (continued)*

Mnemonic	Operands	Function	Other Registers Altered
<b>twnl</b>	RA, RB	Trap if (RA) not less than (RB). <i>Extended mnemonic for</i> <b>tw 12,RA,RB</b>	



**twi** TO, RA, IM


```

if ( ((RA) < EXTS(IM) ^ TO0 = 1) ∨
      ((RA) > EXTS(IM) ^ TO1 = 1) ∨
      ((RA) = EXTS(IM) ^ TO2 = 1) ∨
      ((RA) u EXTS(IM) ^ TO3 = 1) ∨
      ((RA) s EXTS(IM) ^ TO4 = 1) )
  SRR0 ← address of twi instruction
  SRR1 ← MSR
  ESR[PTR] ← 1 (other bits cleared)
  MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] ← 90
  PC ← IVPR0:15 || IVOR616:27 || 40
else no operation

```

Register RA is compared with the sign-extended IM field. If any comparison condition selected by the TO field is true, a Trap exception type Program interrupt occurs as follows (see *Program Interrupt* on page 180 for more information on Program interrupts). The contents of the MSR are copied into SRR1 and the address of the **twi** instruction) is placed into SRR0. ESR[PTR] is set to 1 and the other bits ESR bits cleared to indicate the type of exception causing the Program interrupt.

The program counter (PC) is then loaded with the interrupt vector address. The interrupt vector address is formed by concatenating the high halfword of the Interrupt Vector Prefix Register (IVPR), bits 16:27 of the Interrupt Vector Offset Register 6 (IVOR6), and 0b0000.

MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] are set to 0.

Program execution continues at the new address in the PC.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- SRR0 (if trap condition is met)
- SRR1 (if trap condition is met)
- MSR[WE, EE, PR, FP, FE0, FE1, DWE, DS, IS] (if trap condition is met)
- ESR (if trap condition is met)

### Invalid Instruction Forms

- Reserved fields

### Programming Notes


This instruction can be inserted into the execution stream by a debugger to implement breakpoints, and is not typically used by application code.

The enabling of trap debug events may affect the interrupt type caused by the execution of **tw** instruction. Specifically, trap instructions may be enabled to cause Debug interrupts instead of Program interrupts. See *Trap (TRAP) Debug Event* on page 228 for more details.

Table 9-33. Extended Mnemonics for *twi*

Mnemonic	Operands	Function	Other Registers Altered
<b>tweqi</b>	RA, IM	Trap if (RA) equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 4,RA,IM</b>	
<b>twgei</b>	RA, IM	Trap if (RA) greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 12,RA,IM</b>	
<b>twgti</b>	RA, IM	Trap if (RA) greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 8,RA,IM</b>	
<b>twlei</b>	RA, IM	Trap if (RA) less than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 20,RA,IM</b>	
<b>twlgei</b>	RA, IM	Trap if (RA) logically greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 5,RA,IM</b>	
<b>twlgti</b>	RA, IM	Trap if (RA) logically greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 1,RA,IM</b>	
<b>twllei</b>	RA, IM	Trap if (RA) logically less than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 6,RA,IM</b>	
<b>twllti</b>	RA, IM	Trap if (RA) logically less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 2,RA,IM</b>	
<b>twlngi</b>	RA, IM	Trap if (RA) logically not greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 6,RA,IM</b>	
<b>twlnli</b>	RA, IM	Trap if (RA) logically not less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 5,RA,IM</b>	
<b>twlti</b>	RA, IM	Trap if (RA) less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 16,RA,IM</b>	
<b>twnei</b>	RA, IM	Trap if (RA) not equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 24,RA,IM</b>	
<b>twngi</b>	RA, IM	Trap if (RA) not greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 20,RA,IM</b>	
<b>twnli</b>	RA, IM	Trap if (RA) not less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 12,RA,IM</b>	

**wrtee**            RS



$$\text{MSR}[\text{EE}] \leftarrow (\text{RS})_{16}$$

MSR[EE] is set to the value specified by bit 16 of register RS.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

**Registers Altered**

- MSR[EE]

**Invalid Instruction Forms:**

- Reserved fields

**Programming Notes**

Execution of this instruction is privileged.

This instruction is typically used as part of a code sequence which can provide the equivalent of an atomic read-modify-write of the MSR, as follows:

```
mfmsr Rn    #save EE in Rn[16]
wrteei 0    #Turn off EE (leaving other bits unchanged)
•          #Code with EE disabled
•
•
wrtee Rn    #restore EE without affecting any MSR changes that occurred in the disabled code
```


## wrteei

Write External Enable Immediate  
PPC440x5 CPU Core User's Manual



Preliminary

**wrteei**            E



MSR[EE] ← E

MSR[EE] is set to the value specified by the E field.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

### Registers Altered

- MSR[EE]

### Invalid Instruction Forms:

- Reserved fields

### Programming Notes


Execution of this instruction is privileged.

This instruction is typically used as part of a code sequence which can provide the equivalent of an atomic read-modify-write of the MSR, as follows:

```
mfmsr Rn    #save EE in Rn[16]
wrteei 0    #Turn off EE (leaving other bits unchanged)
•          #Code with EE disabled
•
•
wrtee Rn    #restore EE without affecting any MSR changes that occurred in the disabled code
```



**xor** RA, RS, RB R<sub>c</sub>=0  
**xor.** RA, RS, RB R<sub>c</sub>=1



$$(RA) \leftarrow (RS) \oplus (RB)$$

The contents of register RS are XORed with the contents of register RB; the result is placed into register RA.

**Registers Altered**

- RA
- CR[CR0] if R<sub>c</sub> contains 1

## xori


XOR Immediate

PPC440x5 CPU Core User's Manual



Preliminary

**xori** RA, RS, IM



$$(RA) \leftarrow (RS) \oplus ({}^{16}0 \parallel IM)$$


The IM field is extended to 32 bits by concatenating 16 0-bits on the left. The contents of register RS are XORed with the extended IM field; the result is placed into register RA.

### Registers Altered

- RA



**xoris** RA, RS, IM



$$(RA) \leftarrow (RS) \oplus (IM \parallel 16_0)$$

The IM field is extended to 32 bits by concatenating 16 0-bits on the right. The contents of register RS are XORed with the extended IM field; the result is placed into register RA.

**Registers Altered**

- RA





## 10. Register Summary

This chapter provides an alphabetical listing of and bit definitions for the registers contained in the PPC440x5 core.

The registers, of five types, are grouped into several functional categories according to the processor functions with which they are associated. More information about the registers and register categories is provided in *Section 2.2 Registers* on page 47, and in the chapters describing the processor functions with which each register category is associated.

### 10.1 Register Categories

*Table 10-1* summarizes the register categories and the registers contained in each category. Italicized register names are implementation-specific. All other registers are defined by the Book-E Enhanced PowerPC Architecture.

*Table 10-1. Register Categories*

Register Category	Register(s)	Model and Access	Type	Page
<b>Branch Control</b>	CR	User	CR	66
	CTR	User	SPR	66
	LR	User	SPR	65
<b>Cache Control</b>	<i>DNV0–DNV3</i>	Supervisor	SPR	95
	<i>DTV0–DTV3</i>	Supervisor	SPR	95
	<i>DVLIM</i>	Supervisor	SPR	97
	<i>INV0–INV3</i>	Supervisor	SPR	95
	<i>ITV0–ITV3</i>	Supervisor	SPR	95
	<i>IVLIM</i>	Supervisor	SPR	97
<b>Cache Debug</b>	<i>DCDBTRH, DCDBTRL</i>	Supervisor, read-only	SPR	123
	<i>ICDBDR, ICDBTRH, ICDBTRL</i>	Supervisor, read-only	SPR	109
<b>Debug</b>	DAC1–DAC2	Supervisor	SPR	240
	DBCR0–DBCR2	Supervisor	SPR	233
	DBDR	Supervisor	SPR	241
	DBSR	Supervisor	SPR	238
	DVC1–DVC2	Supervisor	SPR	240
	IAC1–IAC4	Supervisor	SPR	239
<b>Device Control</b>	Implemented outside core	Supervisor	DCR	52
<b>Integer Processing</b>	GPR0–GPR31	User	GPR	69
	XER	User	SPR	70

Table 10-1. Register Categories

Register Category	Register(s)	Model and Access	Type	Page
Interrupt Processing	CSRR0–CSRR1	Supervisor	SPR	162
	DEAR	Supervisor	SPR	164
	ESR	Supervisor	SPR	166
	IVOR0–IVOR15	Supervisor	SPR	164
	IVPR	Supervisor	SPR	165
	MCSR	Supervisor	SPR	168
	MCSRR0–MCSRR1	Supervisor	SPR	163
	SRR0–SRR1	Supervisor	SPR	161
Processor Control	CCR0	Supervisor	SPR	105
	CCR1	Supervisor	SPR	105
	MSR	Supervisor	MSR	159
	PIR, PVR	Supervisor, read-only	SPR	73
	RSTCFG	Supervisor, read-only	SPR	77
	SPRG0–SPRG3	Supervisor	SPR	73
	SPRG4–SPRG7	User, read-only; Supervisor	SPR	73
	USPRG0	User	SPR	73
Storage Control	MMUCR	Supervisor	SPR	143
	PID	Supervisor	SPR	146
Timer	DEC	Supervisor	SPR	205
	DECAR	Supervisor, write-only	SPR	205
	TBL, TBU	User read, Supervisor write	SPR	204
	TCR	Supervisor	SPR	209
	TSR	Supervisor	SPR	210

Table 10-2 *Special Purpose Registers Sorted by SPR Number* on page 443, lists the Special Purpose Registers (SPRs) in order by SPR number (SPRN). The table provides mnemonics, names, SPRN, model (user or supervisor), and access. All SPR numbers not listed are reserved, and should be neither read nor written.

Note that three registers, DBSR, MCSR, and TSR, are indicated as having the access type of *read/clear*. These three registers are *status* registers, and as such behave differently than other SPRs when written. The term “read/clear” does not mean that these registers are automatically cleared upon being read. Rather, the “clear” refers to their behavior when being written. Instead of simply overwriting the SPR with the data in the source GPR, the status SPR is updated by zeroing those bit positions corresponding to 1 values in the source GPR, with those bit positions corresponding to 0 values in the source GPR being left unchanged. In this fashion, it is possible for software to read one of these status SPRs, and then write to it using the same data which was read. Any bits which were read as 1 will then be cleared, and any bits which were not yet set at the time the SPR was read will be left unchanged. If any of these previously clear bits happen to be set between the time the SPR is read and when it is written, then when the SPR is later read again, software will observe any newly set bit[s]. If it were not for this behavior, then software could erroneously clear bits which it had not yet observed as having been set, and overlook the occurrence of certain exceptions.

Table 10-2. Special Purpose Registers Sorted by SPR Number

Mnemonic	Register Name	SPRN	Model	Access
XER	Integer Exception Register	0x001	User	Read/Write
LR	Link Register	0x008	User	Read/Write
CTR	Count Register	0x009	User	Read/Write
DEC	Decrementer	0x016	Supervisor	Read/Write
SRR0	Save/Restore Register 0	0x01A	Supervisor	Read/Write
SRR1	Save/Restore Register 1	0x01B	Supervisor	Read/Write
PID	Process ID	0x030	Supervisor	Read/Write
DECAR	Decrementer Auto-Reload	0x036	Supervisor	Write-only
CSRR0	Critical Save/Restore Register 0	0x03A	Supervisor	Read/Write
CSRR1	Critical Save/Restore Register 1	0x03B	Supervisor	Read/Write
DEAR	Data Exception Address Register	0x03D	Supervisor	Read/Write
ESR	Exception Syndrome Register	0x03E	Supervisor	Read/Write
IVPR	Interrupt Vector Prefix Register	0x03F	Supervisor	Read/Write
USPRG0	User Special Purpose Register General 0	0x100	User	Read/Write
SPRG4	Special Purpose Register General 4	0x104	User	Read-only
SPRG5	Special Purpose Register General 5	0x105	User	Read-only
SPRG6	Special Purpose Register General 6	0x106	User	Read-only
SPRG7	Special Purpose Register General 7	0x107	User	Read-only
TBL	Time Base Lower	0x10C	User	Read-only
TBU	Time Base Upper	0x10D	User	Read-only
SPRG0	Special Purpose Register General 0	0x110	Supervisor	Read/Write
SPRG1	Special Purpose Register General 1	0x111	Supervisor	Read/Write
SPRG2	Special Purpose Register General 2	0x112	Supervisor	Read/Write
SPRG3	Special Purpose Register General 3	0x113	Supervisor	Read/Write
SPRG4	Special Purpose Register General 4	0x114	Supervisor	Write-only
SPRG5	Special Purpose Register General 5	0x115	Supervisor	Write-only
SPRG6	Special Purpose Register General 6	0x116	Supervisor	Write-only
SPRG7	Special Purpose Register General 7	0x117	Supervisor	Write-only
TBL	Time Base Lower	0x11C	Supervisor	Write-only
TBU	Time Base Upper	0x11D	Supervisor	Write-only
PIR	Processor ID Register	0x11E	Supervisor	Read-only
PVR	Processor Version Register	0x11F	Supervisor	Read-only
DBSR	Debug Status Register	0x130	Supervisor	Read/Clear
DBCR0	Debug Control Register 0	0x134	Supervisor	Read/Write
DBCR1	Debug Control Register 1	0x135	Supervisor	Read/Write
DBCR2	Debug Control Register 2	0x136	Supervisor	Read/Write

Table 10-2. Special Purpose Registers Sorted by SPR Number

Mnemonic	Register Name	SPRN	Model	Access
IAC1	Instruction Address Compare 1	0x138	Supervisor	Read/Write
IAC2	Instruction Address Compare 2	0x139	Supervisor	Read/Write
IAC3	Instruction Address Compare 3	0x13A	Supervisor	Read/Write
IAC4	Instruction Address Compare 4	0x13B	Supervisor	Read/Write
DAC1	Data Address Compare 1	0x13C	Supervisor	Read/Write
DAC2	Data Address Compare 2	0x13D	Supervisor	Read/Write
DVC1	Data Value Compare 1	0x13E	Supervisor	Read/Write
DVC2	Data Value Compare 2	0x13F	Supervisor	Read/Write
TSR	Timer Status Register	0x150	Supervisor	Read/Clear
TCR	Timer Control Register	0x154	Supervisor	Read/Write
IVOR0	Interrupt Vector Offset Register 0	0x190	Supervisor	Read/Write
IVOR1	Interrupt Vector Offset Register 1	0x191	Supervisor	Read/Write
IVOR2	Interrupt Vector Offset Register 2	0x192	Supervisor	Read/Write
IVOR3	Interrupt Vector Offset Register 3	0x193	Supervisor	Read/Write
IVOR4	Interrupt Vector Offset Register 4	0x194	Supervisor	Read/Write
IVOR5	Interrupt Vector Offset Register 5	0x195	Supervisor	Read/Write
IVOR6	Interrupt Vector Offset Register 6	0x196	Supervisor	Read/Write
IVOR7	Interrupt Vector Offset Register 7	0x197	Supervisor	Read/Write
IVOR8	Interrupt Vector Offset Register 8	0x198	Supervisor	Read/Write
IVOR9	Interrupt Vector Offset Register 9	0x199	Supervisor	Read/Write
IVOR10	Interrupt Vector Offset Register 10	0x19A	Supervisor	Read/Write
IVOR11	Interrupt Vector Offset Register 11	0x19B	Supervisor	Read/Write
IVOR12	Interrupt Vector Offset Register 12	0x19C	Supervisor	Read/Write
IVOR13	Interrupt Vector Offset Register 13	0x19D	Supervisor	Read/Write
IVOR14	Interrupt Vector Offset Register 14	0x19E	Supervisor	Read/Write
IVOR15	Interrupt Vector Offset Register 15	0x19F	Supervisor	Read/Write
MCSRR0	Machine Check Save Restore Register 0	0x23A	Supervisor	Read/Write
MCSRR1	Machine Check Save Restore Register 1	0x23B	Supervisor	Read/Write
MCSR	Machine Check Status Register	0x23C	Supervisor	Read/Write
INV0	Instruction Cache Normal Victim 0	0x370	Supervisor	Read/Write
INV1	Instruction Cache Normal Victim 1	0x371	Supervisor	Read/Write
INV2	Instruction Cache Normal Victim 2	0x372	Supervisor	Read/Write
INV3	Instruction Cache Normal Victim 3	0x373	Supervisor	Read/Write
ITV0	Instruction Cache Transient Victim 0	0x374	Supervisor	Read/Write
ITV1	Instruction Cache Transient Victim 1	0x375	Supervisor	Read/Write
ITV2	Instruction Cache Transient Victim 2	0x376	Supervisor	Read/Write
ITV3	Instruction Cache Transient Victim 3	0x377	Supervisor	Read/Write

Table 10-2. Special Purpose Registers Sorted by SPR Number

Mnemonic	Register Name	SPRN	Model	Access
CCR1	Core Configuration Register 1	0x378	Supervisor	Read/Write
DNV0	Data Cache Normal Victim 0	0x390	Supervisor	Read/Write
DNV1	Data Cache Normal Victim 1	0x391	Supervisor	Read/Write
DNV2	Data Cache Normal Victim 2	0x392	Supervisor	Read/Write
DNV3	Data Cache Normal Victim 3	0x393	Supervisor	Read/Write
DTV0	Data Cache Transient Victim 0	0x394	Supervisor	Read/Write
DTV1	Data Cache Transient Victim 1	0x395	Supervisor	Read/Write
DTV2	Data Cache Transient Victim 2	0x396	Supervisor	Read/Write
DTV3	Data Cache Transient Victim 3	0x397	Supervisor	Read/Write
DVLIM	Data Cache Victim Limit	0x398	Supervisor	Read/Write
IVLIM	Instruction Cache Victim Limit	0x399	Supervisor	Read/Write
RSTCFG	Reset Configuration	0x39B	Supervisor	Read-only
DCDBTRL	Data Cache Debug Tag Register Low	0x39C	Supervisor	Read-only
DCDBTRH	Data Cache Debug Tag Register High	0x39D	Supervisor	Read-only
ICDBTRL	Instruction Cache Debug Tag Register Low	0x39E	Supervisor	Read-only
ICDBTRH	Instruction Cache Debug Tag Register High	0x39F	Supervisor	Read-only
MMUCR	Memory Management Unit Control Register	0x3B2	Supervisor	Read/Write
CCR0	Core Configuration Register 0	0x3B3	Supervisor	Read/Write
ICDBDR	Instruction Cache Debug Data Register	0x3D3	Supervisor	Read-only
DBDR	Debug Data Register	0x3F3	Supervisor	Read/Write

## 10.2 Reserved Fields

For all registers with fields marked as reserved, the reserved fields should be written as *zero* and read as *undefined*. That is, when writing to a reserved field, write a zero to that field. When reading from a reserved field, ignore that field.

The recommended coding practice is to perform the initial write to a register with reserved fields as described in the preceding paragraph, and to perform all subsequent writes to the register using a read-modify-write strategy: read the register, alter desired fields with logical instructions, and then write the register.

## 10.3 Device Control Registers

Device Control Registers (DCRs), which are architecturally outside of the processor core, are accessed using the **mfdcr** and **mtdcr** instructions. DCRs are used to control, configure, and hold status for various functional units that are not part of the RISC processor core. Although the PPC440x5 core does not contain DCRs, the **mfdcr** and **mtdcr** instructions are provided.

---

The **mfocr** and **mtocr** instructions are privileged, for all DCR numbers. Therefore, all DCR accesses are privileged. All DCR numbers are reserved, and should be neither read nor written, unless they are part of a Core+ASIC implementation.

## 10.4 Alphabetical Register Listing

The following pages list the registers available in the PPC440x5 core. For each register, the following information is supplied:

- Register mnemonic and name
- Cross reference to detailed register information
- Register type (if SPR); the types of the other registers are the same as the register names (CR, GPR, MSR)
- Register number (address)
- Register programming model (user or supervisor) and access (read-clear, read-only, read/write (R/W), write-only)
- A diagram illustrating the register fields (all register fields have mnemonics, unless there is only one field)
- A table describing the register fields, giving field mnemonics, field bit locations, field names, and the functions associated with the various field values

# CCR0

Core Configuration Register 0  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x3B3 Supervisor R/W

See *Core Configuration Register 0 (CCR0)* on page 105.




Figure 10-1. Core Configuration Register 0 (CCR0)

0		Reserved	
1	PRE	Parity Recoverability Enable 0 Semi-recoverable parity mode enabled for data cache 1 Fully recoverable parity mode enabled for data cache	Must be set to 1 to guarantee full recoverability from MMU and data cache parity errors.
2:3		Reserved	
4	CRPE	Cache Read Parity Enable 0 Disable parity information reads 1 Enable parity information reads	When enabled, execution of <b>icread</b> , <b>dcread</b> , or <b>tlbre</b> loads parity information into the ICDBTRH, DCDBTRL, or target GPR, respectively.
5:9		Reserved	
10	DSTG	Disable Store Gathering 0 Enabled; stores to contiguous addresses may be gathered into a single transfer 1 Disabled; all stores to memory will be performed independently	See <i>Store Gathering</i> on page 116.
11	DAPUIB	Disable APU Instruction Broadcast 0 Enabled. 1 Disabled; instructions not broadcast to APU for decoding	This mechanism is provided as a means of reducing power consumption when an auxiliary processor is not attached and/or is not being used. See <i>Initialization</i> on page 83.
12:15		Reserved	
16	DTB	Disable Trace Broadcast 0 Enabled. 1 Disabled; no trace information is broadcast.	This mechanism is provided as a means of reducing power consumption when instruction tracing is not needed. See <i>Initialization</i> on page 83.
17	GICBT	Guaranteed Instruction Cache Block Touch 0 <b>icbt</b> may be abandoned without having filled cache line if instruction pipeline stalls. 1 <b>icbt</b> is guaranteed to fill cache line even if instruction pipeline stalls.	See <i>icbt Operation</i> on page 108.
18	GDCBT	Guaranteed Data Cache Block Touch 0 <b>dcbt/dcbtst</b> may be abandoned without having filled cache line if load/store pipeline stalls. 1 <b>dcbt/dcbtst</b> are guaranteed to fill cache line even if load/store pipeline stalls.	See <i>Data Cache Control and Debug</i> on page 121.
19:22		Reserved	





## Preliminary

23	FLSTA	Force Load/Store Alignment 0 No Alignment exception on integer storage access instructions, regardless of alignment 1 An alignment exception occurs on integer storage access instructions if data address is not on an operand boundary.	See <i>Load and Store Alignment</i> on page 114.
24:27		Reserved	
28:29	ICSLC	Instruction Cache Speculative Line Count	Number of additional lines (0–3) to fill on instruction fetch miss. See <i>Speculative Prefetch Mechanism</i> on page 102.
30:31	ICSLT	Instruction Cache Speculative Line Threshold	Number of doublewords that must have already been filled in order that the current speculative line fill is <i>not</i> abandoned on a redirection of the instruction stream. See <i>Speculative Prefetch Mechanism</i> on page 102.

# CCR1

Core Configuration Register 1  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x378 Supervisor R/W

See *Core Configuration Register 1 (CCR1)* on page 107.




Figure 10-2. Core Configuration Register 1 (CCR1)

0:7	ICDPEI	Instruction Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded when the instruction cache is filled. Each of the 8 bits corresponds to one of the instruction words in the line.
8:9	ICTPEI	Instruction Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the instruction cache.
10:11	DCTPEI	Data Cache Tag Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the data cache.
12	DCDPEI	Data Cache Data Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the data field in the data cache.
13	DCUPEI	Data Cache U-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bit recorded for the U fields in the data cache.
14	DCMPEI	Data Cache Modified-bit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the modified (dirty) field in the data cache.
15	FCOM	Force Cache Operation Miss 0 normal operation 1 cache ops appear to miss the cache	Force <b>icbt</b> , <b>dcbt</b> , <b>dcbtst</b> , <b>dcbst</b> , <b>dcbf</b> , <b>dcbi</b> , and <b>dcbz</b> to appear to miss the caches. The intended use is with <b>icbt</b> and <b>dcbt</b> only, which will fill a duplicate line and allow testing of multi-hit parity errors. See <i>Section 4.2.4.7 Simulating Instruction Cache Parity Errors for Software Testing</i> on page 111 and <i>Figure 4.3.3.7</i> on page 126.
16:19	MMUPEI	Memory Management Unit Parity Error Insert 0 record even parity (normal) 1 record odd parity (simulate parity error)	Controls inversion of parity bits recorded for the tag field in the MMU.
20	FFF	Force Full-line Flush 0 flush only as much data as necessary. 1 always flush entire cache lines	When flushing 32-byte (8-word) lines from the data cache, normal operation is to write nothing, a double word, quad word, or the entire 8-word block to the memory as required by the dirty bits. This bit ensures that none or all dirty bits are set so that either nothing or the entire 8-word block is written to memory when flushing a line from the data cache. Refer to <i>Section 4.3.1.4 Line Flush Operations</i> on page 117.
21:23		Reserved	



Preliminary

24	TCS	Timer Clock Select 0 CPU timer advances by one at each rising edge of the CPU input clock (CPMC440CLOCK). 1 CPU timer advances by one for each rising edge of the CPU timer clock (CPMC440TIMERLOCK). When TCS = 1, CPU timer clock input can toggle at up to half of the CPU clock frequency.
25:31		Reserved

# CR

## User Read/Write

See *Condition Register (CR)* on page 66.




Figure 10-3. Condition Register (CR)

0:3	CR0	Condition Register Field 0
4:7	CR1	Condition Register Field 1
8:11	CR2	Condition Register Field 2
12:15	CR3	Condition Register Field 3
16:19	CR4	Condition Register Field 4
20:23	CR5	Condition Register Field 5
24:27	CR6	Condition Register Field 6
28:31	CR7	Condition Register Field 7



Preliminary

**SPR 0x03A Supervisor R/W**

See *Critical Save/Restore Register 0 (CSRR0)* on page 162.



*Figure 10-4. Critical Save/Restore Register 0 (CSRR0)*

0:29		Return address for critical interrupts
30:31		Reserved

# CSRR1

Critical Save/Restore Register 1  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x03B Supervisor R/W

See *Critical Save/Restore Register 1 (CSRR1)* on page 162.




Figure 10-5. Critical Save/Restore Register 1 (CSRR1)

0:31	Copy of the MSR when a critical interrupt is taken
------	--

Preliminary

**SPR 0x009 User R/W**

See *Count Register (CTR)* on page 66.



*Figure 10-6. Count Register (CTR)*

0:31		Count	Used as count for branch conditional with decrement instructions, or as target address for <b>bcctr</b> instructions
------	--	-------	--

## DAC1–DAC2

Data Address Compare Registers  
PPC440x5 CPU Core User's Manual




Preliminary

### SPR 0x13C–0x13D Supervisor R/W

See *Data Address Compare Registers (DAC1–DAC2)* on page 240.



*Figure 10-7. Data Address Compare Registers (DAC1–DAC2)*





**SPR 0x134 Supervisor R/W**


 See *Debug Control Register 0 (DBCRO)* on page 233.


Figure 10-8. Debug Control Register 0 (DBCRO)

0	EDM	External Debug Mode 0 Disable external debug mode. 1 Enable external debug mode.	
1	IDM	Internal Debug Mode 0 Disable internal debug mode. 1 Enable internal debug mode.	
2:3	RST	Reset 00 No action 01 Core reset 10 Chip reset 11 System reset  <b>Attention:</b> Writing 01, 10, or 11 to this field causes a processor reset to occur.	
4	ICMP	Instruction Completion Debug Event 0 Disable instruction completion debug event. 1 Enable instruction completion debug event.	Instruction completions do not cause instruction completion debug events if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
5	BRT	Branch Taken Debug Event 0 Disable branch taken debug event. 1 Enable branch taken debug event.	Taken branches do not cause branch taken debug events if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
6	IRPT	Interrupt Debug Event 0 Disable interrupt debug event. 1 Enable interrupt debug event.	Critical interrupts do not cause interrupt debug events in internal debug mode, unless also in external debug mode or debug wait mode.
7	TRAP	Trap Debug Event 0 Disable trap debug event. 1 Enable trap debug event.	
8	IAC1	Instruction Address Compare (IAC) 1 Debug Event 0 Disable IAC 1 debug event. 1 Enable IAC 1 debug event.	
9	IAC2	IAC 2 Debug Event 0 Disable IAC 2 debug event. 1 Enable IAC 2 debug event.	
10	IAC3	IAC 3 Debug Event 0 Disable IAC 3 debug event. 1 Enable IAC 3 debug event.	
11	IAC4	IAC 4 Debug Event 0 Disable IAC 4 debug event. 1 Enable IAC 4 debug event.	

## DBCR0 (cont.)

Debug Control Register 0

PPC440x5 CPU Core User's Manual



Preliminary

12	DAC1R	Data Address Compare (DAC) 1 Read Debug Event 0 Disable DAC 1 read debug event. 1 Enable DAC 1 read debug event.	
13	DAC1W	DAC 1 Write Debug Event 0 Disable DAC 1 write debug event. 1 Enable DAC 1 write debug event.	
14	DAC2R	DAC 2 Read Debug Event 0 Disable DAC 2 read debug event. 1 Enable DAC 2 read debug event.	
15	DAC2W	DAC 2 Write Debug Event 0 Disable DAC 2 write debug event. 1 Enable DAC 2 write debug event.	
16	RET	Return Debug Event 0 Disable return ( <b>rfi/rfci/rfmci</b> ) debug event. 1 Enable return ( <b>rfi/rfci/rfmci</b> ) debug event.	<b>rfci/rfmci</b> does not cause a return debug event if MSR[DE] = 0 in internal debug mode, unless also in external debug mode or debug wait mode.
17:30		Reserved	
31	FT	Freeze timers on debug event 0 Timers are not frozen. 1 Freeze timers if a DBSR field associated with a debug event is set.	

**SPR 0x135 Supervisor R/W**


 See *Debug Control Register 1 (DBCR1)* on page 234.


Figure 10-9. Debug Control Register 1 (DBCR1)

0:1	IAC1US	Instruction Address Compare (IAC) 1 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
2:3	IAC1ER	IAC 1 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
4:5	IAC2US	IAC 2 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
6:7	IAC2ER	IAC 2 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
8:9	IAC12M	IAC 1/2 Mode 00 Exact match 01 Reserved 10 Range inclusive 11 Range exclusive	Match if address[0:29] = IAC 1/2[0:29]; two independent compares  Match if IAC1 ≤ address < IAC2 Match if address < IAC1 OR address ≥ IAC2
10:14		Reserved	
15	IAC12AT	IAC 1/2 Auto-Toggle Enable 0 Disable IAC 1/2 auto-toggle 1 Enable IAC 1/2 auto-toggle	
16:17	IAC3US	IAC 3 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
18:19	IAC3ER	IAC 3 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	

## DBCR1 (cont.)

Debug Control Register 1

PPC440x5 CPU Core User's Manual



Preliminary

20:21	IAC4US	IAC 4 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
22:23	IAC4ER	IAC 4 Effective/Real 00 Effective (MSR[IS] = don't care) 01 Reserved 10 Virtual (MSR[IS] = 0) 11 Virtual (MSR[IS] = 1)	
24:25	IAC34M	IAC 3/4 Mode 00 Exact match 01 Reserved 10 Range inclusive 11 Range exclusive	Match if address[0:29] = IAC 3/4[0:29]; two independent compares  Match if $IAC3 \leq \text{address} < IAC4$ Match if $\text{address} < IAC3$ OR $\text{address} \geq IAC4$
26:30		Reserved	
31	IAC34AT	IAC3/4 Auto-Toggle Enable 0 Disable IAC 3/4 auto-toggle 1 Enable IAC 3/4 auto-toggle	

**SPR 0x136 Supervisor R/W**


 See *Debug Control Register 2 (DBCR2)* on page 237.


Figure 10-10. Debug Control Register 2 (DBCR2)

0:1	DAC1US	Data Address Compare (DAC) 1 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
2:3	DAC1ER	DAC 1 Effective/Real 00 Effective (MSR[DS] = don't care) 01 Reserved 10 Virtual (MSR[DS] = 0) 11 Virtual (MSR[DS] = 1)	
4:5	DAC2US	DAC 2 User/Supervisor 00 Both 01 Reserved 10 Supervisor only (MSR[PR] = 0) 11 User only (MSR[PR] = 1)	
6:7	DAC2ER	DAC 2 Effective/Real 00 Effective (MSR[DS] = don't care) 01 Reserved 10 Virtual (MSR[DS] = 0) 11 Virtual (MSR[DS] = 1)	
8:9	DAC12M	DAC 1/2 Mode 00 Exact match 01 Address bit mask 10 Range inclusive 11 Range exclusive	Match if address[0:31] = DAC 1/2[0:31]; two independent compares Match if address = DAC1; only compare bits corresponding to 1 bits in DAC2 Match if DAC1 ≤ address < DAC2 Match if address < DAC1 OR address ≥ DAC2
10	DAC12A	DAC 1/2 Asynchronous 0 Debug interrupt caused by DAC1/2 exception will be synchronous 1 Debug interrupt caused by DAC1/2 exception will be asynchronous	
11		Reserved	
12:13	DVC1M	Data Value Compare (DVC) 1 Mode 00 Reserved 01 AND all bytes enabled by DVC1BE 10 OR all bytes enabled by DVC1BE 11 AND-OR pairs of bytes enabled by DVC1BE	(0 AND 1) OR (2 AND 3)

## DBCR2 (cont.)

Debug Control Register 2

PPC440x5 CPU Core User's Manual



Preliminary

14:15	DVC2M	DVC 2 Mode 00 Reserved 01 AND all bytes enabled by DVC2BE 10 OR all bytes enabled by DVC2BE 11 AND-OR pairs of bytes enabled by DVC2BE (0 AND 1) OR (2 AND 3)
16:19		Reserved
20:23	DVC1BE	DVC 1 Byte Enables 0:3
24:27		Reserved
28:31	DVC2BE	DVC 2 Byte Enables 0:3


Preliminary

**SPR 0x3F3 Supervisor R/W**

See *Debug Data Register (DBDR)* on page 241.



*Figure 10-11. Debug Data Register (DBDR)*



**SPR 0x130 Supervisor Read/Clear**

See *Debug Status Register (DBSR)* on page 238.




Figure 10-12. Debug Status Register (DBSR)

0	IDE	Imprecise Debug Event 0 Debug event occurred while MSR[DE] = 1 1 Debug event occurred while MSR[DE] = 0	For synchronous debug events in internal debug mode, this field indicates whether the corresponding Debug interrupt occurs precisely or imprecisely
1	UDE	Unconditional Debug Event 0 Event didn't occur 1 Event occurred	
2:3	MRR	Most Recent Reset 00 No reset has occurred since this field was last cleared by software. 01 Core reset 10 Chip reset 11 System reset	This field is set upon any processor reset to a value indicating the type of reset.
4	ICMP	Instruction Completion Debug Event 0 Event didn't occur 1 Event occurred	
5	BRT	Branch Taken Debug Event 0 Event didn't occur 1 Event occurred	
6	IRPT	Interrupt Debug Event 0 Event didn't occur 1 Event occurred	
7	TRAP	Trap Debug Event 0 Event didn't occur 1 Event occurred	
8	IAC1	IAC 1 Debug Event 0 Event didn't occur 1 Event occurred	
9	IAC2	IAC 2 Debug Event 0 Event didn't occur 1 Event occurred	
10	IAC3	IAC 3 Debug Event 0 Event didn't occur 1 Event occurred	
11	IAC4	IAC 4 Debug Event 0 Event didn't occur 1 Event occurred	
12	DAC1R	DAC 1 Read Debug Event 0 Event didn't occur 1 Event occurred	





Preliminary

13	DAC1W	DAC 1 Write Debug Event 0 Event didn't occur 1 Event occurred
14	DAC2R	DAC 2 Read Debug Event 0 Event didn't occur 1 Event occurred
15	DAC2W	DAC 2 Write Debug Event 0 Event didn't occur 1 Event occurred
16	RET	Return Debug Event 0 Event didn't occur 1 Event occurred
17:29		Reserved
30	IAC12ATS	IAC 1/2 Auto-Toggle Status 0 Range is not reversed from value specified in DBCR1[IAC12M] 1 Range is reversed from value specified in DBCR1[IAC12M]
31	IAC34ATS	IAC 3/4 Auto-Toggle Status 0 Range is not reversed from value specified in DBCR1[IAC34M] 1 Range is reversed from value specified in DBCR1[IAC34M]

# DCDBTRH

Data Cache Debug Tag Register High  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x39D Supervisor Read-Only

See *dcread Operation* on page 123.




Figure 10-13. Data Cache Debug Tag Register High (DCDBTRH)

0:23	TRA	Tag Real Address	Bits 0:23 of the lower 32 bits of the 36-bit real address associated with the cache line read by <b>dcread</b> .
24	V	Cache Line Valid 0 Cache line is not valid. 1 Cache line is valid.	The valid indicator for the cache line read by <b>dcread</b> .
25:27		Reserved	
28:31	TERA	Tag Extended Real Address	Upper 4 bits of the 36-bit real address associated with the cache line read by <b>dcread</b> .

**SPR 0x39C Supervisor Read-Only**


 See *dcread Operation* on page 123.


Figure 10-14. Data Cache Debug Tag Register Low (DCDBTRL)

0:12		Reserved	
13	UPAR	U bit parity	The parity for the U0-U3 bits in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
14:15	TPAR	Tag parity	The parity for the tag bits in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
16:19	DPAR	Data parity	The parity <i>check values</i> for the data bytes in the word read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
20:23	MPAR	Modified (dirty) parity	The parity for the modified (dirty) indicators for each of the four doublewords in the cache line read by <b>dcread</b> if CCR0[CRPE] = 1, otherwise 0.
24:27	D	Dirty Indicators	The “dirty” (modified) indicators for each of the four doublewords in the cache line read by <b>dcread</b> .
28	U0	U0 Storage Attribute	The U0 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
29	U1	U1 Storage Attribute	The U1 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
30	U2	U2 Storage Attribute	The U2 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .
31	U3	U3 Storage Attribute	The U3 storage attribute for the memory page associated with this cache line read by <b>dcread</b> .

## DEAR


Data Exception Address Register  
PPC440x5 CPU Core User's Manual



Preliminary

### SPR 0x03D Supervisor R/W

See *Data Exception Address Register (DEAR)* on page 164.




*Figure 10-15. Data Exception Address Register (DEAR)*


0:31	Address of data exception for Data Storage, Alignment, and Data TLB Error interrupts
------	--

**SPR 0x016 Supervisor R/W**

See *Decrementer (DEC)* on page 205.



*Figure 10-16. Decrementer (DEC)*



# DECAR

Decrementer Auto-Reload  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x036 Supervisor Write-Only

See *Decrementer (DEC)* on page 205.





Figure 10-17. Decrementer Auto-Reload (DECAR)

0:31	Decrementer auto-reload value	Copied to DEC at next time base clock when DEC = 1 and auto-reload is enabled (TCR[ARE] = 1).
------	-------------------------------	---

Preliminary

**SPR 0x390–0x393 Supervisor R/W**

 See *Cache Line Replacement Policy* on page 94.

*Figure 10-18. Data Cache Normal Victim Registers (DNV0–DNV3)*

0:7	VNDXA	Victim Index A (for cache lines with EA[25:26] = 0b00)	For all victim index fields, the number of bits used to select the cache way for replacement depends on the implemented cache size. See <i>Table 4-5</i> on page 117. for more information.
8:15	VNDXB	Victim Index B (for cache lines with EA[25:26] = 0b01)	
16:23	VNDXC	Victim Index C (for cache lines with EA[25:26] = 0b10)	
24:31	VNDXD	Victim Index D (for cache lines with EA[25:26] = 0b11)	

# DTV0–DTV3

Data Cache Transient Victim 0–3  
PPC440x5 CPU Core User’s Manual



Preliminary

## SPR 0x394–0x397 Supervisor R/W

See *Cache Line Replacement Policy* on page 94.




Figure 10-19. Data Cache Transient Victim Registers (DTV0–DTV3)

0:7	VNDXA	Victim Index A (for cache lines with EA[25:26] = 0b00)	For all victim index fields, the number of bits used to select the cache way for replacement depends on the implemented cache size. See <i>Table 4-3</i> on page 96. for more information.
8:15	VNDXB	Victim Index B (for cache lines with EA[25:26] = 0b01)	
16:23	VNDXC	Victim Index C (for cache lines with EA[25:26] = 0b10)	
24:31	VNDXD	Victim Index D (for cache lines with EA[25:26] = 0b11)	







Preliminary

**SPR 0x13E–0x13F Supervisor R/W**

See *Data Value Compare Registers (DVC1–DVC2)* on page 240.



*Figure 10-20. Data Value Compare Registers (DVC1–DVC2)*



# DVLIM



## SPR 0x398 Supervisor R/W

See *Cache Locking and Transient Mechanism* on page 96.




Figure 10-21. Data Cache Victim Limit (DVLIM)

0:1		Reserved	
2:9	TFLOOR	Transient Floor	The number of bits in the TFLOOR field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.
10:12		Reserved	
13:20	TCEILING	Transient Ceiling	The number of bits in the TCEILING field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.
21:23		Reserved	
24:31	NFLOOR	Normal Floor	The number of bits in the NFLOOR field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.

**SPR 0x03E Supervisor R/W**

See *Exception Syndrome Register (ESR)* on page 166.




Figure 10-22. Exception Syndrome Register (ESR)

0	MCI	Machine Check—Instruction Fetch Exception 0 Instruction Machine Check exception did not occur. 1 Instruction Machine Check exception occurred.	This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.
1:3		Reserved	
4	PIL	Program Interrupt—Illegal Instruction Exception 0 Illegal Instruction exception did not occur. 1 Illegal Instruction exception occurred.	
5	PPR	Program Interrupt—Privileged Instruction Exception 0 Privileged Instruction exception did not occur. 1 Privileged Instruction exception occurred.	
6	PTR	Program Interrupt—Trap Exception 0 Trap exception did not occur. 1 Trap exception occurred.	
7	FP	Floating Point Operation 0 Exception was not caused by a floating point instruction. 1 Exception was caused by a floating point instruction.	
8	ST	Store Operation 0 Exception was not caused by a store-type storage access or cache management instruction. 1 Exception was caused by a store-type storage access or cache management instruction.	
9		Reserved	
10:11	DLK	Data Storage Interrupt—Locking Exception 00 Locking exception did not occur. 01 Locking exception was caused by <b>dcbf</b> . 10 Locking exception was caused by <b>icbi</b> . 11 Reserved	
12	AP	AP Operation 0 Exception was not caused by an auxiliary processor instruction. 1 Exception was caused by an auxiliary processor instruction.	
13	PUO	Program Interrupt—Unimplemented Operation Exception 0 Unimplemented Operation exception did not occur. 1 Unimplemented Operation exception occurred.	

## ESR (cont.)

Exception Syndrome Register

PPC440x5 CPU Core User's Manual




Preliminary

14	BO	Byte Ordering Exception 0 Byte Ordering exception did not occur. 1 Byte Ordering exception occurred.	
15	PIE	Program Interrupt—Imprecise Exception 0 Exception occurred precisely; SRR0 contains the address of the instruction that caused the exception. 1 Exception occurred imprecisely; SRR0 contains the address of an instruction after the one which caused the exception.	This field is only set for a Floating-Point Enabled exception type Program interrupt, and then only when the interrupt occurs imprecisely due to MSR[FE0,FE1] being set to a non-zero value when an attached floating-point unit is already signaling the Floating-Point Enabled exception (that is, FPSCR[FEX] is already 1).
16:26		Reserved	
27	PCRE	Program Interrupt—Condition Register Enable 0 Instruction which caused the exception is not a floating-point CR-updating instruction. 1 Instruction which caused the exception is a floating-point CR-updating instruction.	This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.  This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.
28	PCMP	Program Interrupt—Compare 0 Instruction which caused the exception is not a floating-point compare type instruction 1 Instruction which caused the exception is a floating-point compare type instruction.	This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.  This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.
29:31	PCRF	Program Interrupt—Condition Register Field If ESR[PCRE]=1, this field indicates which CR field was to be updated by the floating-point instruction which caused the exception.	This is an implementation-dependent field of the ESR and is not part of the PowerPC Book-E Architecture.  This field is only defined for a Floating-Point Enabled exception type Program interrupt, and then only when ESR[PIE] is 0.

Preliminary

**User R/W**

See *General Purpose Registers (GPRs)* on page 69.




*Figure 10-23. General Purpose Registers (R0-R31)*

0:31	General Purpose Register data
------	-------------------------------



Preliminary

**SPR 0x3D3 Supervisor Read-Only**

 See *icread Operation* on page 109.

*Figure 10-25. Instruction Cache Debug Data Register (ICDBDR)*

0:31	Instruction machine code from instruction cache
------	---

# ICDBTRH

Instruction Cache Debug Tag Register High  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x39F Supervisor Read-Only

See *icread Operation* on page 109.





Figure 10-26. Instruction Cache Debug Tag Register High (ICDBTRH)

0:23		Tag Effective Address	Bits 0:23 of the 32-bit effective address associated with the cache line read by <b>icread</b> .
24	V	Cache Line Valid 0 Cache line is not valid. 1 Cache line is valid.	The valid indicator for the cache line read by <b>icread</b> .
25:26	TPAR	Tag Parity	The parity bits for the address tag for the cache line read by <b>icread</b> , if CCR0[CRPE] is set.
27	DAPAR	Instruction Data parity	The parity bit for the instruction word at the 32-bit effective address specified in the <b>icread</b> instruction, if CCR0[CRPE] is set.
28:31		Reserved	



Preliminary

**SPR 0x39E Supervisor Read-Only**

 See *icread Operation* on page 109.

*Figure 10-27. Instruction Cache Debug Tag Register Low (ICDBTRL)*

0:21		Reserved	
22	TS	Translation Space	The address space portion of the virtual address associated with the cache line read by <b>icread</b> .
23	TD	Translation ID (TID) Disable 0 TID enable 1 TID disable	TID Disable field for the memory page associated with the cache line read by <b>icread</b> .
24:31	TID	Translation ID	TID field portion of the virtual address associated with the cache line read by <b>icread</b> .

## INV0–INV3

Instruction Cache Normal Victim 0–3  
PPC440x5 CPU Core User's Manual



Preliminary

### SPR 0x370–0x373 Supervisor R/W

See *Cache Line Replacement Policy* on page 94.





Figure 10-28. Instruction Cache Normal Victim Registers (INV0–INV3)

0:7	VNDXA	Victim Index A (for cache lines with EA[25:26] = 0b00)	For all victim index fields, the number of bits used to select the cache way for replacement depends on the implemented cache size. See <i>Table 4-3</i> on page 96. for more information.
8:15	VNDXB	Victim Index B (for cache lines with EA[25:26] = 0b01)	
16:23	VNDXC	Victim Index C (for cache lines with EA[25:26] = 0b10)	
24:31	VNDXD	Victim Index D (for cache lines with EA[25:26] = 0b11)	

**SPR 0x374–0x377 Supervisor R/W**

 See *Cache Line Replacement Policy* on page 94.

**Figure 10-29. Instruction Cache Transient Victim Registers (ITV0–ITV3)**

0:7	VNDXA	Victim Index A (for cache lines with EA[25:26] = 0b00)	For all victim index fields, the number of bits used to select the cache way for replacement depends on the implemented cache size. See <i>Table 4-3</i> on page 96. for more information.
8:15	VNDXB	Victim Index B (for cache lines with EA[25:26] = 0b01)	
16:23	VNDXC	Victim Index C (for cache lines with EA[25:26] = 0b10)	
24:31	VNDXD	Victim Index D (for cache lines with EA[25:26] = 0b11)	

# IVLIM

## SPR 0x399 Supervisor R/W

See *Cache Locking and Transient Mechanism* on page 96.





Figure 10-30. Instruction Cache Victim Limit (IVLIM)

0:1		Reserved	
2:9	TFLOOR	Transient Floor	The number of bits in the TFLOOR field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.
10:12		Reserved	
13:20	TCEILING	Transient Ceiling	The number of bits in the TCEILING field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.
21:23		Reserved	
24:31	NFLOOR	Normal Floor	The number of bits in the NFLOOR field varies, depending on the implemented cache size. See <i>Table 4-3</i> on page 96 for more information.

Preliminary

**SPR 0x190–0x19F Supervisor R/W**

See *Interrupt Vector Offset Registers (IVOR0–IVOR15)* on page 164.



*Figure 10-31. Interrupt Vector Offset Registers (IVOR0–IVOR15)*

0:15		Reserved
16:27	IVO	Interrupt Vector Offset
28:31		Reserved

*Table 10-3. Interrupt Types Associated with each IVOR*

IVOR	Interrupt Type
IVOR0	Critical Input
IVOR1	Machine Check
IVOR2	Data Storage
IVOR3	Instruction Storage
IVOR4	External Input
IVOR5	Alignment
IVOR6	Program
IVOR7	Floating Point Unavailable
IVOR8	System Call
IVOR9	Auxiliary Processor Unavailable
IVOR10	Decrementer
IVOR11	Fixed Interval Timer
IVOR12	Watchdog Timer
IVOR13	Data TLB Error
IVOR14	Instruction TLB Error
IVOR15	Debug

# IVPR

Interrupt Vector Prefix Register  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x03F Supervisor R/W

See *Interrupt Vector Prefix Register (IVPR)* on page 165.





Figure 10-32. *Interrupt Vector Prefix Register (IVPR)*

0:15	IVP	Interrupt Vector Prefix
16:31		Reserved

Preliminary

**SPR 0x008 User R/W**

See *Link Register (LR)* on page 65.



*Figure 10-33. Link Register (LR)*



# MCSR

Machine Check Status Register  
 PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x23C Supervisor Read/Clear

See *Machine Check Status Register (MCSR)* on page 168.





Figure 10-34. Machine Check Status Register (MCSR)

0	MCS	Machine Check Summary 0 No async machine check exception pending 1 Async machine check exception pending	Set when a machine check exception occurs that is handled in the asynchronous fashion. One of MCSR bits 1:7 will be set simultaneously to indicate the exception type. When MSR[ME] and this bit are both set, Machine Check interrupt is taken.
1	IB	Instruction PLB Error 0 Exception not caused by Instruction Read PLB interrupt request (IRQ) 1 Exception caused by Instruction Read PLB interrupt request (IRQ)	
2	DRB	Data Read PLB Error 0 Exception not caused by Data Read PLB interrupt request (IRQ) 1 Exception caused by Data Read PLB interrupt request (IRQ)	
3	DWB	Data Write PLB Error 0 Exception not caused by Data Write PLB interrupt request (IRQ) 1 Exception caused by Data Write PLB interrupt request (IRQ)	
4	TLBP	Translation Lookaside Buffer Parity Error 0 Exception not caused by TLB parity error 1 Exception caused by TLB parity error	
5	ICP	Instruction Cache Parity Error 0 Exception not caused by I-cache parity error 1 Exception caused by I-cache parity error	
6	DCSP	Data Cache Search Parity Error 0 Exception not caused by DCU Search parity error 1 Exception caused by DCU Search parity error	Set if and only if the DCU parity error was discovered during a DCU Search operation. See <i>Data Cache Parity Operations</i> on page 125.
7	DCFP	Data Cache Flush Parity Error 0 Exception not caused by DCU Flush parity error 1 Exception caused by DCU Flush parity error	Set if and only if the DCU parity error was discovered during a DCU Flush operation. See <i>Data Cache Parity Operations</i> on page 125.
8	IMPE	Imprecise Machine Check Exception 0 No imprecise machine check exception occurred. 1 Imprecise machine check exception occurred.	Set if a machine check exception occurs that sets MCSR[MCS] (or would if it were not already set) and MSR[ME] = 0.
9:31		Reserved	



Preliminary

**SPR 0x23A Supervisor R/W**

 See *Machine Check Save/Restore Register 0 (MCSRR0)* on page 163.

*Figure 10-35. Machine Check Save/Restore Register 0 (MCSRR0)*

0:29		Return address for machine check interrupts
30:31		Reserved

# MCSRR1

Machine Check Save/Restore Register 1  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x23B Supervisor R/W

See *Machine Check Save/Restore Register 1 (MCSRR1)* on page 163.




Figure 10-36. Machine Check Save/Restore Register 1 (MCSRR1)

0:31	Copy of the MSR at the time of a machine check interrupt.
------	---

**SPR 0x3B2 Supervisor R/W**


 See *Memory Management Unit Control Register (MMUCR)* on page 143.


Figure 10-37. Memory Management Unit Control Register (MMUCR)

0:6		Reserved	
7	SWOA	Store Without Allocate 0 Cacheable store misses allocate a line in the data cache. 1 Cacheable store misses do not allocate a line in the data cache.	If MMUCR[U2SWOAE] = 1, this field is ignored.
8		Reserved	
9	U1TE	U1 Transient Enable 0 Disable U1 storage attribute as transient storage attribute. 1 Enable U1 storage attribute as transient storage attribute.	
10	U2SWOAE	U2 Store without Allocate Enable 0 Disable U2 storage attribute control of store without allocate. 1 Enable U2 storage attribute control of store without allocate.	If MMUCR[U2SWOAE] = 1, the U2 storage attribute overrides MMUCR[SWOA].
11		Reserved	
12	DULXE	Data Cache Unlock Exception Enable 0 Data cache unlock exception is disabled. 1 Data cache unlock exception is enabled.	<b>dcbf</b> in user mode causes a Cache Locking exception type Data Storage interrupt when MMUCR[DULXE] is 1.
13	IULXE	Instruction Cache Unlock Exception Enable 0 Instruction cache unlock exception is disabled. 1 Instruction cache unlock exception is enabled.	<b>icbi</b> in user mode causes a Cache Locking exception type Data Storage interrupt when MMUCR[IULXE] is 1.
14		Reserved	
15	STS	Search Translation Space	Specifies the value of the translation space (TS) field for <b>tlbsx[.]</b>
16:23		Reserved	
24:31	STID	Search Translation ID	Specifies the value of the TID field for the <b>tlbsx[.]</b> ; also used to transfer a TLB entry's TID value for <b>tlbre</b> and <b>tlbwe</b> .

# MSR



## Supervisor R/W

See *Machine State Register (MSR)* on page 159.




Figure 10-38. Machine State Register (MSR)

0:12		Reserved	
13	WE	Wait State Enable 0 The processor is not in the wait state. 1 The processor is in the wait state.	If MSR[WE] = 1, the processor remains in the wait state until an interrupt is taken, a reset occurs, or an external debug tool clears WE.
14	CE	Critical Interrupt Enable 0 Critical Input and Watchdog Timer interrupts are disabled. 1 Critical Input and Watchdog Timer interrupts are enabled.	
15		Reserved	
16	EE	External Interrupt Enable 0 External Input, Decrementer, and Fixed Interval Timer interrupts are disabled. 1 External Input, Decrementer, and Fixed Interval Timer interrupts are enabled.	
17	PR	Problem State 0 Supervisor state (privileged instructions can be executed) 1 Problem state (privileged instructions can not be executed)	
18	FP	Floating Point Available 0 The processor cannot execute floating-point instructions 1 The processor can execute floating-point instructions	
19	ME	Machine Check Enable 0 Machine Check interrupts are disabled 1 Machine Check interrupts are enabled.	
20	FE0	Floating-point exception mode 0 0 If MSR[FE1] = 0, ignore exceptions mode; if MSR[FE1] = 1, imprecise nonrecoverable mode 1 If MSR[FE1] = 0, imprecise recoverable mode; if MSR[FE1] = 1, precise mode	
21	DWE	Debug Wait Enable 0 Disable debug wait mode. 1 Enable debug wait mode.	
22	DE	Debug interrupt Enable 0 Debug interrupts are disabled. 1 Debug interrupts are enabled.	



23	FE1	Floating-point exception mode 1 0 If MSR[FE0] = 0, ignore exceptions mode; if MSR[FE0] = 1, imprecise recoverable mode 1 If MSR[FE0] = 0, imprecise non-recoverable mode; if MSR[FE0] = 1, precise mode
24:25		Reserved
26	IS	Instruction Address Space 0 All instruction storage accesses are directed to address space 0 (TS = 0 in the relevant TLB entry). 1 All instruction storage accesses are directed to address space 1 (TS = 1 in the relevant TLB entry).
27	DS	Data Address Space 0 All data storage accesses are directed to address space 0 (TS = 0 in the relevant TLB entry). 1 All data storage accesses are directed to address space 1 (TS = 1 in the relevant TLB entry).
28:31		Reserved

# PID

Process ID

PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x030 Supervisor R/W

See *Process ID (PID)* on page 146.





Figure 10-39. Process ID (PID)

0:23		Reserved
24:31	PID	Process ID

Preliminary

**SPR 0x11E Supervisor Read-Only**

 See *Processor Identification Register (PIR)* on page 74.

*Figure 10-40. Processor Identification Register (PIR)*

0:27		Reserved
28:31	PIN	Processor Identification Number (PIN)

# PVR

## SPR 0x11F Supervisor Read-Only

See *Processor Version Register (PVR)* on page 73.




Figure 10-41. Processor Version Register (PVR)

0:11	OWN	Owner Identifier	Identifies the owner of a core.
12:31	PVN	Processor Version Number	Implementation-specific value identifying the specific version and use of a processor core within a chip.



**SPR 39B Supervisor Read-Only**


 See *Reset Configuration (RSTCFG)* on page 77.


Figure 10-42. Reset Configuration

0:15		Reserved	
16	U0	U0 Storage Attribute 0 U0 storage attribute is disabled 1 U0 storage attribute is enabled	See Table 5-1 on page 131.
17	U1	U1 Storage Attribute 0 U1 storage attribute is disabled 1 U1 storage attribute is enabled	See Table 5-1 on page 131.
18	U2	U2 Storage Attribute 0 U2 storage attribute is disabled 1 U2 storage attribute is enabled	See Table 5-1 on page 131.
19	U3	U3 Storage Attribute 0 U3 storage attribute is disabled 1 U3 storage attribute is enabled	See Table 5-1 on page 131.
20:23		Reserved	
24	E	E Storage Attribute 0 Accesses to the page are big endian. 1 Accesses to the page are little endian.	
25:27		Reserved	
28:31	ERPEN	Extended Real Page Number	This TLB field is prepended to the translated address to form a 36-bit real address. See Table 5.4 Address Translation on page 136 and Table 5-3 Page Size and Real Address Formation on page 137.

## SPRG0–SPRG7

Special Purpose Register General  
PPC440x5 CPU Core User's Manual



Preliminary

**SPR 0x104–0x107 (User/Supervisor Read-Only); SPR 0x110–0x113 (Supervisor R/W);  
SPR 0x114–0x117 (Supervisor Write-Only)**

See *Special Purpose Registers General (USPRG0, SPRG0–SPRG7)* on page 73.




*Figure 10-43. Special Purpose Registers General (SPRG0–SPRG7)*

0:31	General data	Software value; no hardware usage.
------	--------------	------------------------------------

Preliminary

**SPR 0x01A Supervisor R/W**

See *Save/Restore Register 0 (SRR0)* on page 161.



*Figure 10-44. Save/Restore Register 0 (SRR0)*

0:29		Return address for non-critical interrupts
30:31		Reserved

# SRR1

Save/Restore Register 1  
PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x01B Supervisor R/W

See *Save/Restore Register 1 (SRR1)* on page 161.





Figure 10-45. *Save/Restore Register 1 (SRR1)*

0:31	Copy of the MSR at the time of a non-critical interrupt.
------	--

**SPR 0x10C (User/Supervisor Read-Only); SPR 0x11C (Supervisor Write-Only)**

See *Time Base* on page 204.



*Figure 10-46. Time Base Lower (TBL)*

0:31	Time Base Lower	Low-order 32 bits of time base.
------	-----------------	---------------------------------

# TBU

Time Base Upper


PPC440x5 CPU Core User's Manual




Preliminary

**SPR 0x10D (User/Supervisor Read-Only); SPR 0x11D (Supervisor Write-Only)**


See *Time Base* on page 204.



*Figure 10-47. Time Base Upper (TBU)*



**SPR 0x154 Supervisor R/W**

 See *Timer Control Register (TCR)* on page 209.

 Figure 10-48. *Timer Control Register (TCR)*

0:1	WP	Watchdog Timer Period 00 $2^{21}$ time base clocks 01 $2^{25}$ time base clocks 10 $2^{29}$ time base clocks 11 $2^{33}$ time base clocks	
2:3	WRC	Watchdog Timer Reset Control 00 No Watchdog Timer reset will occur. 01 Core reset 10 Chip reset 11 System reset	TCR[WRC] resets to 0b00. Type of reset to cause upon Watchdog Timer exception with TSR[ENW,WIS]=0b11. This field can be set by software, but cannot be cleared by software, except by a software-induced reset.
4	WIE	Watchdog Timer Interrupt Enable 0 Disable Watchdog Timer interrupt. 1 Enable Watchdog Timer interrupt.	
5	DIE	Decrementer Interrupt Enable 0 Disable Decrementer interrupt. 1 Enable Decrementer interrupt.	
6:7	FP	Fixed Interval Timer (FIT) Period 00 $2^{13}$ time base clocks 01 $2^{17}$ time base clocks 10 $2^{21}$ time base clocks 11 $2^{25}$ time base clocks	
8	FIE	FIT Interrupt Enable 0 Disable Fixed Interval Timer interrupt. 1 Enable Fixed Interval Timer interrupt.	
9	ARE	Auto-Reload Enable 0 Disable auto reload. 1 Enable auto reload.	TCR[ARE] resets to 0b0.
10:31		Reserved	

# TSR

Timer Status Register

PPC440x5 CPU Core User's Manual



Preliminary

## SPR 0x150 Supervisor Read/Clear

See *Timer Status Register (TSR)* on page 210.




Figure 10-49. *Timer Status Register (TSR)*


0	ENW	Enable Next Watchdog Timer Exception 0 Action on next Watchdog Timer exception is to set TSR[ENW] = 1. 1 Action on next Watchdog Timer exception is governed by TSR[WIS].
1	WIS	Watchdog Timer Interrupt Status 0 Watchdog Timer exception has not occurred. 1 Watchdog Timer exception has occurred.
2:3	WRS	Watchdog Timer Reset Status 00 No Watchdog Timer reset has occurred. 01 Core reset was forced by Watchdog Timer. 10 Chip reset was forced by Watchdog Timer. 11 System reset was forced by Watchdog Timer.
4	DIS	Decrementer Interrupt Status 0 Decrementer exception has not occurred. 1 Decrementer exception has occurred.
5	FIS	Fixed Interval Timer (FIT) Interrupt Status 0 Fixed Interval Timer exception has not occurred. 1 Fixed Interval Timer exception has occurred.
6:31		Reserved



Preliminary

**SPR 0x100 (User R/W)**

See *Special Purpose Registers General (USPRG0, SPRG0–SPRG7)* on page 73.



*Figure 10-50. User Special Purpose Register General (USPRG0)*

0:31	General data	Software value; no hardware usage.
------	--------------	------------------------------------

# XER

## SPR 0x001 User R/W

See *Integer Exception Register (XER)* on page 70.




Figure 10-51. Integer Exception Register (XER)

0	SO	Summary Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by <b>mtspr</b> or by integer or auxiliary processor instructions with the [o] option; can be <i>reset</i> by <b>mtspr</b> or by <b>mcrxr</b> .
1	OV	Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by <b>mtspr</b> or by integer or allocated instructions with the [o] option; can be <i>reset</i> by <b>mtspr</b> , by <b>mcrxr</b> , or by integer or allocated instructions with the [o] option.
2	CA	Carry 0 Carry has not occurred. 1 Carry has occurred.	Can be <i>set</i> by <b>mtspr</b> or by certain integer arithmetic and shift instructions; can be <i>reset</i> by <b>mtspr</b> , by <b>mcrxr</b> , or by certain integer arithmetic and shift instructions.
3:24		Reserved	
25:31	TBC	Transfer Byte Count	Used as a byte count by <b>lswx</b> and <b>stswx</b> ; written by <b>dlimzb[.]</b> and by <b>mtspr</b> .

## Appendix A. Instruction Summary

This appendix describes the various instruction formats, and lists all of the PPC440x5 instructions summarized alphabetically and by opcode.

*Appendix A.1* on page 507 illustrates the PPC440x5 instruction forms (allowed arrangements of fields within instructions).

*Appendix A.2* on page 512 lists all PPC440x5 mnemonics, including extended mnemonics. A short functional description is included for each mnemonic.

*Appendix A.3* on page 543 identifies those opcodes which are allocated by PowerPC Book-E for implementation-dependent usage, including auxiliary processors.

*Appendix A.4* on page 543 identifies those opcodes which are identified by PowerPC Book-E as “preserved” for compatibility with previous versions of the architecture.

*Appendix A.5* on page 544 identifies those opcodes which are “reserved” for use by future versions of the architecture.

*Appendix A.6* on page 544, lists all instructions implemented within the PPC440x5 core, sorted by primary and secondary opcodes. Extended mnemonics are not included in the opcode list, but allocated, preserved, and reserved-nop opcodes are included.

### A.1 Instruction Formats

Instructions are four bytes long. Instruction addresses are always word-aligned.

Instruction bits 0 through 5 always contain the primary opcode. Many instructions have an extended opcode in another field. Remaining instruction bits contain additional fields. All instruction fields belong to one of the following categories:

- Defined

These instructions contain values, such as opcodes, that cannot be altered. The instruction format diagrams specify the values of defined fields.

- Variable

These fields contain operands, such as GPR selectors and immediate values, that can vary from execution to execution. The instruction format diagrams specify the operands in the variable fields.

- Reserved

Bits in reserved fields should be set to 0. In the instruction format diagrams, /, //, or /// indicate reserved fields.

If any bit in a defined field does not contain the expected value, the instruction is illegal and an illegal instruction exception occurs. If any bit in a reserved field does not contain 0, the instruction form is invalid; its result is architecturally undefined. The PPC440x5 core executes all invalid instruction forms without causing an illegal instruction exception.

### A.1.1 Instruction Fields

PPC440x5 instructions contain various combinations of the following fields, as indicated in the instruction format diagrams that follow the field definitions. Numbers, enclosed in parentheses, that follow the field names indicate bit positions; bit fields are indicated by starting and stopping bit positions separated by colons.

AA (30)	Absolute address bit.  0 The immediate field represents an address relative to the current instruction address (CIA). The effective address (EA) of the branch is either the sum of the LI field sign-extended to 32 bits and the branch instruction address, or the sum of the BD field sign-extended to 32 bits and the branch instruction address.  1 The immediate field represents an absolute address. The EA of the branch is either the LI field or the BD field, sign-extended to 32 bits.
BA (11:15)	Specifies a bit in the CR used as a source of a CR-logical instruction.
BB (16:20)	Specifies a bit in the CR used as a source of a CR-logical instruction.
BD (16:29)	An immediate field specifying a 14-bit signed twos complement branch displacement. This field is concatenated on the right with 0b00 and sign-extended to 32 bits.
BF (6:8)	Specifies a field in the CR used as a target in a compare or <b>mcrf</b> instruction.
BFA (11:13)	Specifies a field in the CR used as a source in a <b>mcrf</b> instruction.
BI (11:15)	Specifies a bit in the CR used as a source for the condition of a conditional branch instruction.
BO (6:10)	Specifies options for conditional branch instructions. See <i>Branch Instruction BO Field</i> on page 63.
BT (6:10)	Specifies a bit in the CR used as a target as the result of a CR-Logical instruction.
D (16:31)	Specifies a 16-bit signed two's-complement integer displacement for load/store instructions.
DCRF (11:20)	Specifies a device control register (DCR). This field represents the DCR Number (DCRN) with the upper and lower five bits reversed (that is, DCRF = DCRN[5:9]    DCRN[0:4]).
FXM (12:19)	Field mask used to identify CR fields to be updated by the <b>mtcrf</b> instruction.
IM (16:31)	An immediate field used to specify a 16-bit value (either signed integer or unsigned).
LI (6:29)	An immediate field specifying a 24-bit signed twos complement branch displacement; this field is concatenated on the right with b'00' and sign-extended to 32 bits.
LK (31)	Link bit.  0 Do not update the link register (LR).  1 Update the LR with the address of the next instruction.
MB (21:25)	Mask begin.  Used in rotate-and-mask instructions to specify the beginning bit of a mask.
ME (26:30)	Mask end.  Used in rotate-and-mask instructions to specify the ending bit of a mask.
MO (6:10)	Memory Ordering.  Provides a storage ordering function for storage accesses executing prior to an <b>mbar</b> instruction. MO is ignored and treated as 0 in the PPC440x5 CPU core.
NB (16:20)	Specifies the number of bytes to move in an immediate string load or store.

OPCD (0:5)	Primary opcode. Primary opcodes, in decimal, appear in the instruction format diagrams presented with individual instructions. The OPCD field name does not appear in instruction descriptions.
OE (21)	Enables setting the OV and SO fields in the fixed-point exception register (XER) for extended arithmetic.
RA (11:15)	A GPR used as a source or target.
RB (16:20)	A GPR used as a source.
Rc (31)	Record bit. 0 Do not set the CR. 1 Set the CR to reflect the result of an operation. See <i>Condition Register (CR)</i> on page 66 for a further discussion of how the CR bits are set.
RS (6:10)	A GPR used as a source.
RT (6:10)	A GPR used as a target.
SH (16:20)	Specifies a shift amount.
SPRF (11:20)	Specifies a special purpose register (SPR). This field represents the SPR Number (SPRN) with the upper and lower five bits reversed (that is, SPRF = SPRN[5:9]    SPRN[0:4]).
TO (6:10)	Specifies the conditions on which to trap, as described under <b>tw</b> and <b>twi</b> instructions.
WS (16:20)	Specifies the portion of a TLB entry to be read/written by <b>tlbre/tlbwe</b> .
XO (21:30)	Extended opcode for instructions without an OE field. Extended opcodes, in decimal, appear in the instruction format diagrams presented with individual instructions. The XO field name does not appear in instruction descriptions.
XO (22:30)	Extended opcode for instructions with an OE field. Extended opcodes, in decimal, appear in the instruction format diagrams presented with individual instructions. The XO field name does not appear in instruction descriptions.

### A.1.2 Instruction Format Diagrams

The instruction formats (also called “forms”) illustrated in Figure A-1 through Figure A-9 are valid combinations of instruction fields. *Table A-5* on page 545 indicates which “form” is utilized by each PPC440x5 opcode. Fields indicated by slashes (*/*, *//*, or *///*) are reserved. The figures are adapted from the PowerPC User Instruction Set Architecture.

**A.1.2.1 I-Form**




Figure A-1. I Instruction Format

**A.1.2.2 B-Form**



Figure A-2. B Instruction Format

**A.1.2.3 SC-Form**



Figure A-3. SC Instruction Format

**A.1.2.4 D-Form**




Figure A-4. D Instruction Format

**A.1.2.5 X-Form**

OPCD	RT		RA		RB		XO		Rc
OPCD	RT		RA		RB		XO		/
OPCD	RT		RA		NB		XO		/
OPCD	RT		RA		WS		XO		/
OPCD	RT		///		RB		XO		/
OPCD	RT		///		///		XO		/
OPCD	RS		RA		RB		XO		Rc
OPCD	RS		RA		RB		XO		1
OPCD	RS		RA		RB		XO		/
OPCD	RS		RA		NB		XO		/
OPCD	RS		RA		WS		XO		/
OPCD	RS		RA		SH		XO		Rc
OPCD	RS		RA		///		XO		Rc
OPCD	RS		///		RB		XO		/
OPCD	RS		///		///		XO		/
OPCD	BF	/	L		RA		RB		XO
OPCD	BF	//		BFA	//		///		XO
OPCD	BF	//		///		///			XO
OPCD	BF	//		///		U			XO
OPCD	BF	//		///		///			XO
OPCD	TO			RA		RB			XO
OPCD	BT			///		///			XO
OPCD	MO			///		///			XO
OPCD	///			RA		RB			XO
OPCD	///			///		///			XO
OPCD	///			///	E	//			XO
0	6		11		16		21		31

Figure A-5. X Instruction Format

**A.1.2.6 XL-Form**

OPCD	BT	BA	BB	XO	/		
OPCD	BC	BI	///	XO	LK		
OPCD	BF	//	BFA	//	///	XO	/
OPCD	///	///	///	XO	/		
0	6	11	16	21	31		

Figure A-6. XL Instruction Format

**A.1.2.7 XFX-Form**

OPCD	RT	SPRF	XO	/		
OPCD	RT	DCRF	XO	/		
OPCD	RT	/	FXM	/	XO	/
OPCD	RS	SPRF	XO	/		
OPCD	RS	DCRF	XO	/		
0	6	11	16	21	31	

Figure A-7. XFX Instruction Format

**A.1.2.8 XO-Form**

OPCD	RT	RA	RB	OE	XO	Rc
OPCD	RT	RA	RB	OE	XO	Rc
OPCD	RT	RA	///	/	XO	Rc
0	6	11	16	21	22	31

Figure A-8. XO Instruction Format

**A.1.2.9 M-Form**

OPCD	RS	RA	RB	MB	ME	Rc
OPCD	RS	RA	SH	MB	ME	Rc
0	6	11	16	21	26	31

Figure A-9. M Instruction Format

**A.2 Alphabetical Summary of Implemented Instructions**



Table A-1 summarizes the PPC440x5 instruction set, including required extended mnemonics. All mnemonics are listed alphabetically, without regard to whether the mnemonic is realized in hardware or software. When an instruction supports multiple hardware mnemonics (for example, **b**, **ba**, **bl**, **bla** are all forms of **b**), the instruction is alphabetized under the root form. The hardware instructions are described in detail in Chapter 9, "Instruction Set," which is also alphabetized under the root form. Section 9 also describes the instruction operands and notation.

**Programming Note:** Bit 4 of the BO instruction field provides a hint about the most likely outcome of a conditional branch. (See *Branch Prediction* on page 64 for a detailed description of branch prediction.) Assemblers should set  $BO_4 = 0$  unless a specific reason exists otherwise. In the BO field values specified in Table A-1,  $BO_4 = 0$  has always been assumed. The assembler must enable the programmer to specify branch prediction. To do this, the assembler supports suffixes for the conditional branch mnemonics:

+ Predict branch to be taken.

– Predict branch not to be taken.

For example, **bc** also could be coded as **bc+** or **bc–**, and **bne** also could be coded **bne+** or **bne–**. These alternate codings set  $BO_4 = 1$  only if the requested prediction differs from the standard prediction. See *Branch Prediction* on page 64 for more information.

Table A-1. PPC440x5 Instruction Syntax Summary

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>add</b>	RT, RA, RB	Add (RA) to (RB). Place result in RT.		248
<b>add.</b>			CR[CR0]	
<b>addo</b>			XER[SO, OV]	
<b>addo.</b>			CR[CR0] XER[SO, OV]	
<b>addc</b>	RT, RA, RB	Add (RA) to (RB). Place result in RT. Place carry-out in XER[CA].		249
<b>addc.</b>			CR[CR0]	
<b>addco</b>			XER[SO, OV]	
<b>addco.</b>			CR[CR0] XER[SO, OV]	
<b>adde</b>	RT, RA, RB	Add XER[CA], (RA), (RB). Place result in RT. Place carry-out in XER[CA].		250
<b>adde.</b>			CR[CR0]	
<b>addeo</b>			XER[SO, OV]	
<b>addeo.</b>			CR[CR0] XER[SO, OV]	
<b>addi</b>	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT.		251
<b>addic</b>	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].		252
<b>addic.</b>	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].	CR[CR0]	253
<b>addis</b>	RT, RA, IM	Add (IM    <sup>16</sup> 0) to (RA 0). Place result in RT.		254

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>addme</b>	RT, RA	Add XER[CA], (RA), (-1). Place result in RT. Place carry-out in XER[CA].		255
<b>addme.</b>			CR[CR0]	
<b>addmeo</b>			XER[SO, OV]	
<b>addmeo.</b>			CR[CR0] XER[SO, OV]	
<b>addze</b>	RT, RA	Add XER[CA] to (RA). Place result in RT. Place carry-out in XER[CA].		256
<b>addze.</b>			CR[CR0]	
<b>addzeo</b>			XER[SO, OV]	
<b>addzeo.</b>			CR[CR0] XER[SO, OV]	
<b>and</b>	RA, RS, RB	AND (RS) with (RB). Place result in RA.		257
<b>and.</b>			CR[CR0]	
<b>andc</b>	RA, RS, RB	AND (RS) with $\neg$ (RB). Place result in RA.		258
<b>andc.</b>			CR[CR0]	
<b>andi.</b>	RA, RS, IM	AND (RS) with ( $^{16}0 \parallel$ IM). Place result in RA.	CR[CR0]	259
<b>andis.</b>	RA, RS, IM	AND (RS) with (IM $\parallel$ $^{16}0$ ). Place result in RA.	CR[CR0]	260
<b>b</b>	target	Branch unconditional relative. $LI \leftarrow (\text{target} - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^2_0)$		261
<b>ba</b>		Branch unconditional absolute. $LI \leftarrow \text{target}_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^2_0)$		
<b>bl</b>		Branch unconditional relative. $LI \leftarrow (\text{target} - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^2_0)$	(LR) $\leftarrow$ CIA + 4	
<b>bla</b>		Branch unconditional absolute. $LI \leftarrow \text{target}_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^2_0)$	(LR) $\leftarrow$ CIA + 4	
<b>bc</b>	BO, BI, target	Branch conditional relative. $BD \leftarrow (\text{target} - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^2_0)$	CTR if $BO_2 = 0$	262
<b>bca</b>		Branch conditional absolute. $BD \leftarrow \text{target}_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^2_0)$	CTR if $BO_2 = 0$	
<b>bcl</b>		Branch conditional relative. $BD \leftarrow (\text{target} - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^2_0)$	CTR if $BO_2 = 0$ (LR) $\leftarrow$ CIA + 4	
<b>bcla</b>		Branch conditional absolute. $BD \leftarrow \text{target}_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^2_0)$	CTR if $BO_2 = 0$ (LR) $\leftarrow$ CIA + 4	
<b>bcctr</b>	BO, BI	Branch conditional to address in CTR.	CTR if $BO_2 = 0$	268
<b>bcctrl</b>		Using (CTR) at exit from instruction, $NIA \leftarrow CTR_{0:29} \parallel ^2_0$	CTR if $BO_2 = 0$ (LR) $\leftarrow$ CIA + 4	
<b>bclr</b>	BO, BI	Branch conditional to address in LR.	CTR if $BO_2 = 0$	271
<b>bclrl</b>		Using (LR) at entry to instruction, $NIA \leftarrow LR_{0:29} \parallel ^2_0$	CTR if $BO_2 = 0$ (LR) $\leftarrow$ CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bctr</b>		Branch unconditionally to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 20,0</b>		268
<b>bctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 20,0</b>	(LR) ← CIA + 4	
<b>bdnz</b>	target	Decrement CTR. Branch if CTR ≠ 0 <i>Extended mnemonic for</i> <b>bc 16,0,target</b>		262
<b>bdnza</b>		<i>Extended mnemonic for</i> <b>bca 16,0,target</b>		
<b>bdnzl</b>		<i>Extended mnemonic for</i> <b>bcl 16,0,target</b>	(LR) ← CIA + 4	
<b>bdnzla</b>		<i>Extended mnemonic for</i> <b>bcla 16,0,target</b>	(LR) ← CIA + 4	
<b>bdnzlr</b>		Decrement CTR. Branch if CTR ≠ 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 16,0</b>		271
<b>bdnzlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 16,0</b>	(LR) ← CIA + 4	
<b>bdnzf</b>	cr_bit, target	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 0 <i>Extended mnemonic for</i> <b>bc 0,cr_bit,target</b>		262
<b>bdnzfa</b>		<i>Extended mnemonic for</i> <b>bca 0,cr_bit,target</b>		
<b>bdnzfl</b>		<i>Extended mnemonic for</i> <b>bcl 0,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdnzfla</b>		<i>Extended mnemonic for</i> <b>bcla 0,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdnzflr</b>	cr_bit	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 0,cr_bit</b>		271
<b>bdnzflrl</b>		<i>Extended mnemonic for</i> <b>bclrl 0,cr_bit</b>	(LR) ← CIA + 4	
<b>bdnz8t</b>	cr_bit, target	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 8,cr_bit,target</b>		262
<b>bdnz8ta</b>		<i>Extended mnemonic for</i> <b>bca 8,cr_bit,target</b>		
<b>bdnz8tl</b>		<i>Extended mnemonic for</i> <b>bcl 8,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdnz8tla</b>		<i>Extended mnemonic for</i> <b>bcla 8,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdnz8tlr</b>	cr_bit	Decrement CTR. Branch if CTR ≠ 0 AND CR <sub>cr_bit</sub> = 1 to address in LR. <i>Extended mnemonic for</i> <b>bclr 8,cr_bit</b>		271
<b>bdnz8ttrl</b>		<i>Extended mnemonic for</i> <b>bclrl 8,cr_bit</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bdz</b>	target	Decrement CTR. Branch if CTR = 0 <i>Extended mnemonic for</i> <b>bc 18,0,target</b>		262
<b>bdza</b>		<i>Extended mnemonic for</i> <b>bca 18,0,target</b>		
<b>bdzl</b>		<i>Extended mnemonic for</i> <b>bcl 18,0,target</b>	(LR) ← CIA + 4	
<b>bdzla</b>		<i>Extended mnemonic for</i> <b>bcla 18,0,target</b>	(LR) ← CIA + 4	
<b>bdzlr</b>	cr_bit	Decrement CTR. Branch if CTR = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 18,0</b>		271
<b>bdzlr</b>		<i>Extended mnemonic for</i> <b>bclr 18,0</b>	(LR) ← CIA + 4	
<b>bdzf</b>	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 0 <i>Extended mnemonic for</i> <b>bc 2,cr_bit,target</b>		262
<b>bdzfa</b>		<i>Extended mnemonic for</i> <b>bca 2,cr_bit,target</b>		
<b>bdzfl</b>		<i>Extended mnemonic for</i> <b>bcl 2,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdzfla</b>		<i>Extended mnemonic for</i> <b>bcla 2,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdzflr</b>	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 2,cr_bit</b>		271
<b>bdzflr</b>		<i>Extended mnemonic for</i> <b>bclr 2,cr_bit</b>	(LR) ← CIA + 4	
<b>bdzt</b>	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 10,cr_bit,target</b>		262
<b>bdzta</b>		<i>Extended mnemonic for</i> <b>bca 10,cr_bit,target</b>		
<b>bdztl</b>		<i>Extended mnemonic for</i> <b>bcl 10,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdztla</b>		<i>Extended mnemonic for</i> <b>bcla 10,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bdztlr</b>	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR <sub>cr_bit</sub> = 1 to address in LR. <i>Extended mnemonic for</i> <b>bclr 10,cr_bit</b>		271
<b>bdztlr</b>		<i>Extended mnemonic for</i> <b>bclr 10,cr_bit</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>beq</b>	[cr_field], target	Branch if equal. UseCR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+2,target</b>		262
<b>beqa</b>		<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+2,target</b>		
<b>beql</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+2,target</b>	(LR) ← CIA + 4	
<b>beqla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+2,target</b>	(LR) ← CIA + 4	
<b>beqctr</b>	[cr_field]	Branch if equal to address in CTR. UseCR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+2</b>		268
<b>beqctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+2</b>	(LR) ← CIA + 4	
<b>beqlr</b>	[cr_field]	Branch if equal to address in LR. UseCR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+2</b>		271
<b>beqlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+2</b>	(LR) ← CIA + 4	
<b>bf</b>	cr_bit, target	Branch if CR <sub>cr_bit</sub> = 0. <i>Extended mnemonic for</i> <b>bc 4,cr_bit,target</b>		262
<b>bfa</b>		<i>Extended mnemonic for</i> <b>bca 4,cr_bit,target</b>		
<b>bfl</b>		<i>Extended mnemonic for</i> <b>bcl 4,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bfla</b>		<i>Extended mnemonic for</i> <b>bcla 4,cr_bit,target</b>	(LR) ← CIA + 4	
<b>bfctr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 0 to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 4,cr_bit</b>		268
<b>bfctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,cr_bit</b>	(LR) ← CIA + 4	
<b>bflr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 0 to address in LR. <i>Extended mnemonic for</i> <b>bclr 4,cr_bit</b>		271
<b>bflrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,cr_bit</b>	(LR) ← CIA + 4	
<b>bge</b>	[cr_field], target	Branch if greater than or equal. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+0,target</b>		262
<b>bgea</b>		<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+0,target</b>		
<b>bgel</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+0,target</b>	(LR) ← CIA + 4	
<b>bgeLa</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+0,target</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bgectr</b>	[cr_field]	Branch if greater than or equal to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+0</b>		268
<b>bgectrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+0</b>	(LR) ← CIA + 4	
<b>bgeItr</b>	[cr_field]	Branch if greater than or equal to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+0</b>		271
<b>bgeItrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+0</b>	(LR) ← CIA + 4	
<b>bgt</b>		Branch if greater than. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+1,target</b>		262
<b>bgta</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+1,target</b>		
<b>bgtl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>bgsla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>bgtctr</b>	[cr_field]	Branch if greater than to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+1</b>		268
<b>bgtctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+1</b>	(LR) ← CIA + 4	
<b>bgtItr</b>	[cr_field]	Branch if greater than to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+1</b>		271
<b>bgtItrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+1</b>	(LR) ← CIA + 4	
<b>ble</b>		Branch if less than or equal. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+1,target</b>		262
<b>blea</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+1,target</b>		
<b>blel</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>blela</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>blectr</b>	[cr_field]	Branch if less than or equal to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+1</b>		268
<b>blectrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+1</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>blelr</b>	[cr_field]	Branch if less than or equal to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+1</b>		271
<b>blelrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+1</b>	(LR) ← CIA + 4	
<b>blr</b>		Branch unconditionally to address in LR. <i>Extended mnemonic for</i> <b>bclr 20,0</b>		271
<b>blrl</b>		<i>Extended mnemonic for</i> <b>bclrl 20,0</b>	(LR) ← CIA + 4	
<b>blt</b>		Branch if less than. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+0,target</b>		262
<b>blta</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+0,target</b>		
<b>bltl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+0,target</b>	(LR) ← CIA + 4	
<b>bltla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+0,target</b>	(LR) ← CIA + 4	
<b>bltctr</b>	[cr_field]	Branch if less than to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+0</b>		268
<b>bltctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+0</b>	(LR) ← CIA + 4	
<b>bltir</b>	[cr_field]	Branch if less than to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+0</b>		271
<b>bltlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+0</b>	(LR) ← CIA + 4	
<b>bne</b>		Branch if not equal. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+2,target</b>		262
<b>bnea</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+2,target</b>		
<b>bnel</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+2,target</b>	(LR) ← CIA + 4	
<b>bnela</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+2,target</b>	(LR) ← CIA + 4	
<b>bnectr</b>	[cr_field]	Branch if not equal to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+2</b>		268
<b>bnectrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+2</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bnelr</b>	[cr_field]	Branch if not equal to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+2</b>		271
<b>bnelrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+2</b>	(LR) ← CIA + 4	
<b>bng</b>		Branch if not greater than. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+1,target</b>		
<b>bnga</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+1,target</b>		262
<b>bngl</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>bngla</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+1,target</b>	(LR) ← CIA + 4	
<b>bngctr</b>	[cr_field]	Branch if not greater than to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+1</b>		268
<b>bngctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+1</b>	(LR) ← CIA + 4	
<b>bnglr</b>	[cr_field]	Branch if not greater than to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+1</b>		271
<b>bnglrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+1</b>	(LR) ← CIA + 4	
<b>bnl</b>		Branch if not less than. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+0,target</b>		
<b>bnla</b>	[cr_field], target	<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+0,target</b>		262
<b>bnll</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+0,target</b>	(LR) ← CIA + 4	
<b>bnlla</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+0,target</b>	(LR) ← CIA + 4	
<b>bnlctr</b>	[cr_field]	Branch if not less than to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+0</b>		268
<b>bnlctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+0</b>	(LR) ← CIA + 4	
<b>bnllr</b>	[cr_field]	Branch if not less than to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+0</b>		271
<b>bnllrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+0</b>	(LR) ← CIA + 4	



Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bns</b>	[cr_field], target	Branch if not summary overflow. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+3,target</b>		262
<b>bnsa</b>		<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+3,target</b>		
<b>bnsi</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bnsia</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bnsctr</b>	[cr_field]	Branch if not summary overflow to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+3</b>		268
<b>bnsctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bnslr</b>	[cr_field]	Branch if not summary overflow to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+3</b>		271
<b>bnslrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bnu</b>	[cr_field], target	Branch if not unordered. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 4,4*cr_field+3,target</b>		262
<b>bnua</b>		<i>Extended mnemonic for</i> <b>bca 4,4*cr_field+3,target</b>		
<b>bnul</b>		<i>Extended mnemonic for</i> <b>bcl 4,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bnula</b>		<i>Extended mnemonic for</i> <b>bcla 4,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bnuctr</b>	[cr_field]	Branch if not unordered to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 4,4*cr_field+3</b>		268
<b>bnuctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 4,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bnulr</b>	[cr_field]	Branch if not unordered to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 4,4*cr_field+3</b>		271
<b>bnulrl</b>		<i>Extended mnemonic for</i> <b>bclrl 4,4*cr_field+3</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bso</b>	[cr_field], target	Branch if summary overflow. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+3,target</b>		262
<b>bsoa</b>		<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+3,target</b>		
<b>bsol</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bsola</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bsoctr</b>	[cr_field]	Branch if summary overflow to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+3</b>		268
<b>bsoctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bsolr</b>	[cr_field]	Branch if summary overflow to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+3</b>		271
<b>bsolrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bt</b>	cr_bit, target	Branch if CR <sub>cr_bit</sub> = 1. <i>Extended mnemonic for</i> <b>bc 12,cr_bit,target</b>		262
<b>bta</b>		<i>Extended mnemonic for</i> <b>bca 12,cr_bit,target</b>		
<b>btl</b>		<i>Extended mnemonic for</i> <b>bcl 12,cr_bit,target</b>	(LR) ← CIA + 4	
<b>btla</b>		<i>Extended mnemonic for</i> <b>bcla 12,cr_bit,target</b>	(LR) ← CIA + 4	
<b>btctr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 1 to address in CTR. <i>Extended mnemonic for</i> <b>bcctr 12,cr_bit</b>		268
<b>btctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,cr_bit</b>	(LR) ← CIA + 4	
<b>btlr</b>	cr_bit	Branch if CR <sub>cr_bit</sub> = 1, to address in LR. <i>Extended mnemonic for</i> <b>bclr 12,cr_bit</b>		271
<b>btlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,cr_bit</b>	(LR) ← CIA + 4	
<b>bun</b>	[cr_field], target	Branch if unordered. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bc 12,4*cr_field+3,target</b>		262
<b>buna</b>		<i>Extended mnemonic for</i> <b>bca 12,4*cr_field+3,target</b>		
<b>bunl</b>		<i>Extended mnemonic for</i> <b>bcl 12,4*cr_field+3,target</b>	(LR) ← CIA + 4	
<b>bunla</b>		<i>Extended mnemonic for</i> <b>bcla 12,4*cr_field+3,target</b>	(LR) ← CIA + 4	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>bunctr</b>	[cr_field]	Branch if unordered to address in CTR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bcctr 12,4*cr_field+3</b>		268
<b>bunctrl</b>		<i>Extended mnemonic for</i> <b>bcctrl 12,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>bunlr</b>	[cr_field]	Branch if unordered, to address in LR. Use CR[CR0] if cr_field is omitted. <i>Extended mnemonic for</i> <b>bclr 12,4*cr_field+3</b>		271
<b>bunlrl</b>		<i>Extended mnemonic for</i> <b>bclrl 12,4*cr_field+3</b>	(LR) ← CIA + 4	
<b>clrlwi</b>	RA, RS, n	Clear left immediate. (n < 32) (RA) <sub>0:n-1</sub> ← <sup>n</sup> 0 <i>Extended mnemonic for</i> <b>rlwinm RA,RS,0,n,31</b>		392
<b>clrlwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,0,n,31</b>	CR[CR0]	
<b>clrlslwi</b>	RA, RS, b, n	Clear left and shift left immediate. (n ≤ b < 32) (RA) <sub>b-n:31-n</sub> ← (RS) <sub>b:31</sub> (RA) <sub>32-n:31</sub> ← <sup>n</sup> 0 (RA) <sub>0:b-n-1</sub> ← <sup>b-n</sup> 0 <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,b-n,31-n</b>		392
<b>clrlslwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,b-n,31-n</b>	CR[CR0]	
<b>clrrwi</b>	RA, RS, n	Clear right immediate. (n < 32) (RA) <sub>32-n:31</sub> ← <sup>n</sup> 0 <i>Extended mnemonic for</i> <b>rlwinm RA,RS,0,0,31-n</b>		392
<b>clrrwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,0,0,31-n</b>	CR[CR0]	
<b>cmp</b>	BF, 0, RA, RB	Compare (RA) to (RB), signed. Results in CR[CRn], where n = BF.		275
<b>cmpi</b>	BF, 0, RA, IM	Compare (RA) to EXTS(IM), signed. Results in CR[CRn], where n = BF.		276
<b>cmpl</b>	BF, 0, RA, RB	Compare (RA) to (RB), unsigned. Results in CR[CRn], where n = BF.		277
<b>cmpli</b>	BF, 0, RA, IM	Compare (RA) to ( <sup>16</sup> 0    IM), unsigned. Results in CR[CRn], where n = BF.		278
<b>cmplw</b>	[BF,] RA, RB	Compare Logical Word. Use CR[CR0] if BF is omitted. <i>Extended mnemonic for</i> <b>cmpl BF,0,RA,RB</b>		277
<b>cmplwi</b>	[BF,] RA, IM	Compare Logical Word Immediate. Use CR[CR0] if BF is omitted. <i>Extended mnemonic for</i> <b>cmpli BF,0,RA,IM</b>		278

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>cmpw</b>	[BF,] RA, RB	Compare Word. UseCR[CR0] if BF is omitted. <i>Extended mnemonic for</i> <b>cmp BF,0,RA,RB</b>		275
<b>cmpwi</b>	[BF,] RA, IM	Compare Word Immediate. UseCR[CR0] if BF is omitted. <i>Extended mnemonic for</i> <b>cmpi BF,0,RA,IM</b>		276
<b>cntlzw</b> <b>cntlzw.</b>	RA, RS	Count leading zeros in RS. Place result in RA.	CR[CR0]	279
<b>crand</b>	BT, BA, BB	AND bit (CR <sub>BA</sub> ) with (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		280
<b>crandc</b>	BT, BA, BB	AND bit (CR <sub>BA</sub> ) with $\neg$ (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		281
<b>crclr</b>	bx	Condition register clear. <i>Extended mnemonic for</i> <b>crxor bx,bx,bx</b>		287
<b>creqv</b>	BT, BA, BB	Equivalence of bit CR <sub>BA</sub> with CR <sub>BB</sub> . $CR_{BT} \leftarrow \neg(CR_{BA} \oplus CR_{BB})$		282
<b>crmove</b>	bx, by	Condition register move. <i>Extended mnemonic for</i> <b>cror bx,by,by</b>		285
<b>crnand</b>	BT, BA, BB	NAND bit (CR <sub>BA</sub> ) with (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		283
<b>crnor</b>	BT, BA, BB	NOR bit (CR <sub>BA</sub> ) with (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		284
<b>crnot</b>	bx, by	Condition register not. <i>Extended mnemonic for</i> <b>crnor bx,by,by</b>		284
<b>cror</b>	BT, BA, BB	OR bit (CR <sub>BA</sub> ) with (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		285
<b>crorc</b>	BT, BA, BB	OR bit (CR <sub>BA</sub> ) with $\neg$ (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		286
<b>crset</b>	bx	Condition register set. <i>Extended mnemonic for</i> <b>creqv bx,bx,bx</b>		282
<b>crxor</b>	BT, BA, BB	XOR bit (CR <sub>BA</sub> ) with (CR <sub>BB</sub> ). Place result in CR <sub>BT</sub> .		287
<b>dcba</b>	RA, RB	Treated as a no-op.		288
<b>dcbf</b>	RA, RB	Flush (store, then invalidate) the data cache block which contains the effective address (RA 0) + (RB).		289
<b>dcbi</b>	RA, RB	Invalidate the data cache block which contains the effective address (RA 0) + (RB).		290
<b>dcbst</b>	RA, RB	Store the data cache block which contains the effective address (RA 0) + (RB).		291
<b>dcbt</b>	RA, RB	Load the data cache block which contains the effective address (RA 0) + (RB).		292
<b>dcbtst</b>	RA, RB	Load the data cache block which contains the effective address (RA 0) + (RB).		293
<b>dcbz</b>	RA, RB	Zero the data cache block which contains the effective address (RA 0) + (RB).		294
<b>dccci</b>	RA, RB	Invalidate the data cache array.		295

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>dcread</b>	RT, RA, RB	Read tag and data information from the data cache line selected using effective address bits 17:26. The effective address is calculated by (RA 0) + (RB). Place the data word selected by effective address bits 27:29 in GPR RT; place the tag information in DCDBTRH and DCDBTRL.		296
<b>divw</b>	RT, RA, RB	Divide (RA) by (RB), signed. Place result in RT.	CR[CR0]	298
<b>divw.</b>			XER[SO, OV]	
<b>divwo</b>			CR[CR0]	
<b>divwo.</b>			XER[SO, OV]	
<b>divwu</b>	RT, RA, RB	Divide (RA) by (RB), unsigned. Place result in RT.	CR[CR0]	299
<b>divwu.</b>			XER[SO, OV]	
<b>divwuo</b>			CR[CR0]	
<b>divwuo.</b>			XER[SO, OV]	
<b>dlimzb</b>	RA, RS, RB	$d \leftarrow (RS) \parallel (RB)$ $i, x, y \leftarrow 0$ do while $(x < 8) \wedge (y = 0)$ $x \leftarrow x + 1$ if $d_{i:i+7} = 0$ then $y \leftarrow 1$ else $i \leftarrow i + 8$ $(RA) \leftarrow x$ $XER[TBC] \leftarrow x$ if $Rc = 1$ then $CR[CR0]_3 \leftarrow XER[SO]$ if $y = 1$ then if $x < 5$ then $CR[CR0]_{0:2} \leftarrow 0b010$ else $CR[CR0]_{0:2} \leftarrow 0b100$ else $CR[CR0]_{0:2} \leftarrow 0b001$	XER[TBC], RA	300
<b>dlimzb.</b>			XER[TBC], RA, CR[CR0]	
<b>eqv</b>	RA, RS, RB	Equivalence of (RS) with (RB). $(RA) \leftarrow \neg((RS) \oplus (RB))$		301
<b>eqv.</b>			CR[CR0]	
<b>extlwi</b>	RA, RS, n, b	Extract and left justify immediate. ( $n > 0$ ) $(RA)_{0:n-1} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{n:31} \leftarrow 32-n_0$ Extended mnemonic for <b>rlwinm RA,RS,b,0,n-1</b>		392
<b>extlwi.</b>			Extended mnemonic for <b>rlwinm. RA,RS,b,0,n-1</b> CR[CR0]	
<b>extrwi</b>	RA, RS, n, b	Extract and right justify immediate. ( $n > 0$ ) $(RA)_{32-n:31} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{0:31-n} \leftarrow 32-n_0$ Extended mnemonic for <b>rlwinm RA,RS,b+n,32-n,31</b>		392
<b>extrwi.</b>			Extended mnemonic for <b>rlwinm. RA,RS,b+n,32-n,31</b> CR[CR0]	
<b>extsb</b>	RA, RS	Extend the sign of byte (RS) <sub>24:31</sub> . Place the result in RA.		302
<b>extsb.</b>			CR[CR0]	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>extsh</b>	RA, RS	Extend the sign of halfword (RS) <sub>16:31</sub> .	CR[CR0]	303
<b>extsh.</b>		Place the result in RA.		
<b>icbi</b>	RA, RB	Invalidate the instruction cache block which contains the effective address (RA 0) + (RB).		304
<b>icbt</b>	RA, RB	Load the instruction cache block which contains the effective address (RA 0) + (RB).		302
<b>iccci</b>	RA, RB	Invalidate the instruction cache array.		307
<b>icread</b>	RA, RB	Read tag and data information from the instruction cache line selected using effective address bits 17:26. The effective address is calculated by (RA 0) + (RB). Place the instruction selected by effective address bits 27:29 in ICDBDR; place the tag information in ICDBTRH and ICDBTRL.		308
<b>inslwi</b>	RA, RS, n, b	Insert from left immediate. (n > 0) (RA) <sub>b:b+n-1</sub> ← (RS) <sub>0:n-1</sub> <i>Extended mnemonic for</i> <b>rlwimi RA,RS,32-b,b,b+n-1</b>	CR[CR0]	391
<b>inslwi.</b>		<i>Extended mnemonic for</i> <b>rlwimi. RA,RS,32-b,b,b+n-1</b>		
<b>insrwi</b>	RA, RS, n, b	Insert from right immediate. (n > 0) (RA) <sub>b:b+n-1</sub> ← (RS) <sub>32-n:31</sub> <i>Extended mnemonic for</i> <b>rlwimi RA,RS,32-b-n,b,b+n-1</b>	CR[CR0]	391
<b>insrwi.</b>		<i>Extended mnemonic for</i> <b>rlwimi. RA,RS,32-b-n,b,b+n-1</b>		
<b>isel</b>	RT, RA, RB, CRb	RT ← (RA 0) if CRb = 1, else RT ← (RB)		310
<b>isync</b>		Synchronize execution context by flushing the prefetch queue.		311
<b>la</b>	RT, D(RA)	Load address. (RA ≠ 0) D is an offset from a base address that is assumed to be (RA). (RT) ← (RA) + EXTS(D) <i>Extended mnemonic for</i> <b>addi RT,RA,D</b>		251
<b>lbz</b>	RT, D(RA)	Load byte from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) ← <sup>24</sup> 0    MS(EA,1).		312
<b>lbzu</b>	RT, D(RA)	Load byte from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) ← <sup>24</sup> 0    MS(EA,1). Update the base address, (RA) ← EA.		313
<b>lbzux</b>	RT, RA, RB	Load byte from EA = (RA 0) + (RB) and pad left with zeroes, (RT) ← <sup>24</sup> 0    MS(EA,1). Update the base address, (RA) ← EA.		314
<b>lbzx</b>	RT, RA, RB	Load byte from EA = (RA 0) + (RB) and pad left with zeroes, (RT) ← <sup>24</sup> 0    MS(EA,1).		315
<b>lha</b>	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and sign extend, (RT) ← EXTS(MS(EA,2)).		316
<b>lhau</b>	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and sign extend, (RT) ← EXTS(MS(EA,2)). Update the base address, (RA) ← EA.		317

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>lhaux</b>	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and sign extend, (RT) ← EXTS(MS(EA,2)). Update the base address, (RA) ← EA.		318
<b>lhax</b>	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and sign extend, (RT) ← EXTS(MS(EA,2)).		319
<b>lhbrx</b>	RT, RA, RB	Load halfword from EA = (RA 0) + (RB), then reverse byte order and pad left with zeroes, (RT) ← <sup>16</sup> 0    MS(EA+1,1)    MS(EA,1).		320
<b>lhz</b>	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) ← <sup>16</sup> 0    MS(EA,2).		321
<b>lhzu</b>	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) ← <sup>16</sup> 0    MS(EA,2). Update the base address, (RA) ← EA.		322
<b>lhzux</b>	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) ← <sup>16</sup> 0    MS(EA,2). Update the base address, (RA) ← EA.		323
<b>lhzx</b>	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) ← <sup>16</sup> 0    MS(EA,2).		324
<b>li</b>	RT, IM	Load immediate. (RT) ← EXTS(IM) <i>Extended mnemonic for</i> <b>addi RT,0,value</b>		251
<b>lis</b>	RT, IM	Load immediate shifted. (RT) ← (IM    <sup>16</sup> 0) <i>Extended mnemonic for</i> <b>addis RT,0,value</b>		254
<b>lmw</b>	RT, D(RA)	Load multiple words starting from EA = (RA 0) + EXTS(D). Place into consecutive registers RT through GPR(31). RA is not altered unless RA = GPR(31).		325
<b>lswi</b>	RT, RA, NB	Load consecutive bytes from EA=(RA 0). Number of bytes n=32 if NB=0, else n=NB. Stack bytes into words in CEIL(n/4) consecutive registers starting with RT, to R <sub>FINAL</sub> ← ((RT + CEIL(n/4) – 1) % 32). GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R <sub>FINAL</sub> .		326
<b>lswx</b>	RT, RA, RB	Load consecutive bytes from EA=(RA 0)+(RB). Number of bytes n=XER[TBC]. Stack bytes into words in CEIL(n/4) consecutive registers starting with RT, to R <sub>FINAL</sub> ← ((RT + CEIL(n/4) – 1) % 32). GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R <sub>FINAL</sub> . RB is not altered unless RB = R <sub>FINAL</sub> . If n=0, content of RT is undefined.		328
<b>lwarx</b>	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4). Set the Reservation bit.		330

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>lwbrx</b>	RT, RA, RB	Load word from EA = (RA 0) + (RB) then reverse byte order, (RT) ← MS(EA+3,1)    MS(EA+2,1)    MS(EA+1,1)    MS(EA,1).		331
<b>lwz</b>	RT, D(RA)	Load word from EA = (RA 0) + EXTS(D) and place in RT, (RT) ← MS(EA,4).		332
<b>lwzu</b>	RT, D(RA)	Load word from EA = (RA 0) + EXTS(D) and place in RT, (RT) ← MS(EA,4). Update the base address, (RA) ← EA.		333
<b>lwzux</b>	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4). Update the base address, (RA) ← EA.		334
<b>lwzx</b>	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4).		335
<b>macchw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ (RT) ← temp <sub>1:32</sub>		336
<b>macchw.</b>			CR[CR0]	
<b>macchw0</b>			XER[SO, OV]	
<b>macchw0.</b>			CR[CR0] XER[SO, OV]	
<b>macchwu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ (RT) ← temp <sub>1:32</sub>		339
<b>macchwu.</b>			CR[CR0]	
<b>macchwuo</b>			XER[SO, OV]	
<b>macchwuo.</b>			CR[CR0] XER[SO, OV]	
<b>macchws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ if ((prod <sub>0</sub> = RT <sub>0</sub> ) ∧ (RT <sub>0</sub> ≠ temp <sub>1</sub> )) then (RT) ← (RT <sub>0</sub> ∨ <sup>31</sup> (-RT <sub>0</sub> )) else (RT) ← temp <sub>1:32</sub>		337
<b>macchws.</b>			CR[CR0]	
<b>macchwso</b>			XER[SO, OV]	
<b>macchwso.</b>			CR[CR0] XER[SO, OV]	
<b>macchwsu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ (RT) ← (temp <sub>1:32</sub> ∨ <sup>31</sup> temp <sub>0</sub> )		338
<b>macchwsu.</b>			CR[CR0]	
<b>macchwsuo</b>			XER[SO, OV]	
<b>macchwsuo.</b>			CR[CR0] XER[SO, OV]	
<b>machhw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ (RT) ← temp <sub>1:32</sub>		340
<b>machhw.</b>			CR[CR0]	
<b>machhwo</b>			XER[SO, OV]	
<b>machhwo.</b>			CR[CR0] XER[SO, OV]	
<b>machhwu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ (RT) ← temp <sub>1:32</sub>		343
<b>machhwu.</b>			CR[CR0]	
<b>machhwuo</b>			XER[SO, OV]	
<b>machhwuo.</b>			CR[CR0] XER[SO, OV]	



Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>machws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ if $((prod_0 = RT_0) \wedge (RT_0 \neq temp_1))$ then $(RT) \leftarrow (RT_0 \vee {}^{31}(-RT_0))$ else $(RT) \leftarrow temp_{1:32}$		341
<b>machws.</b>			CR[CR0]	
<b>machwso</b>			XER[SO, OV]	
<b>machwso.</b>			CR[CR0] XER[SO, OV]	
<b>machwsu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee {}^{31}temp_0)$		342
<b>machwsu.</b>			CR[CR0]	
<b>machwsuo</b>			XER[SO, OV]	
<b>machwsuo.</b>			CR[CR0] XER[SO, OV]	
<b>maclhw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		344
<b>maclhw.</b>			CR[CR0]	
<b>maclhwo</b>			XER[SO, OV]	
<b>maclhwo.</b>			CR[CR0] XER[SO, OV]	
<b>maclhwu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		347
<b>maclhwu.</b>			CR[CR0]	
<b>maclhwuo</b>			XER[SO, OV]	
<b>maclhwuo.</b>			CR[CR0] XER[SO, OV]	
<b>maclhws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ if $((prod_0 = RT_0) \wedge (RT_0 \neq temp_1))$ then $(RT) \leftarrow (RT_0 \vee {}^{31}(-RT_0))$ else $(RT) \leftarrow temp_{1:32}$		345
<b>maclhws.</b>			CR[CR0]	
<b>maclhwso</b>			XER[SO, OV]	
<b>maclhwso.</b>			CR[CR0] XER[SO, OV]	
<b>maclhwsu</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee {}^{31}temp_0)$		346
<b>maclhwsu.</b>			CR[CR0]	
<b>maclhwsuo</b>			XER[SO, OV]	
<b>maclhwsuo.</b>			CR[CR0] XER[SO, OV]	
<b>mbar</b>		Storage synchronization. All loads and stores that precede the <b>mbar</b> instruction complete before any loads and stores that follow the instruction access main storage.		348
<b>mcrf</b>	BF, BFA	Move CR field, $(CR[CRn]) \leftarrow (CR[CRm])$ where $m \leftarrow BFA$ and $n \leftarrow BF$		350
<b>mcrxr</b>	BF	Move XER[0:3] into field CRn, where $n \leftarrow BF$ . $CR[CRn] \leftarrow (XER[SO, OV, CA])$ $(XER[SO, OV, CA]) \leftarrow {}^30$		351
<b>mfcrr</b>	RT	Move from CR to RT, $(RT) \leftarrow (CR)$ .		352
<b>mfdcrr</b>	RT, DCRN	Move from DCR to RT, $(RT) \leftarrow (DCR(DCRN))$ .		353
<b>mfmsrr</b>	RT	Move from MSR to RT, $(RT) \leftarrow (MSR)$ .		354

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mfccr0 mfccr1 mfcsrr0 mfcsrr1 mfctr mfdac1 mfdac2 mfdbcr0 mfdbcr1 mfdbcr2 mfdbdr mfdbsr mfdcdbtrh mfdcdbtrl mfdear mfddec mfdnv0 mfdnv1 mfdnv2 mfdnv3 mfdtv0 mfdtv1 mfdtv2 mfdtv3 mfdvc1 mfdvc2 mfdvlim mfer mfiac1 mfiac2 mfiac3 mfiac4 mficdbdr mficdbtrh mficdbtrl mfinv0 mfinv1 mfinv2 mfinv3 mfitv0 mfitv1 mfitv2 mfitv3 mfivlim mfivor0 mfivor1 mfivor2 mfivor3 mfivor4 mfivor5 mfivor6 mfivor7 mfivor8 mfivor9 mfivor10 mfivor11 mfivor12 mfivor13 mfivor14 mfivor15 mfivpr mflr mfmcsr mfmcsrr0 mfmcsrr1 mfmmucr	RT	Move from special purpose register (SPR) SPRN. Extended mnemonic for <b>mf spr RT,SPRN</b> See Table 10-2 Special Purpose Registers Sorted by SPR Number on page 443 for listing of valid SPRN values.		355

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mfpid mfpir mfpvr mfsprg0 mfsprg1 mfsprg2 mfsprg3 mfsprg4 mfsprg5 mfsprg6 mfsprg7 mfsrr0 mfsrr1 mftbl mftbu mftcr mftsr mfsprg0 mfxer		Move from special purpose register (SPR) SPRN. <i>Extended mnemonic for</i> <b>mfspr RT,SPRN</b> See Table 10-2 Special Purpose Registers Sorted by SPR Number on page 443 for listing of valid SPRN values.		
mfspr	RT, SPRN	Move from SPR to RT, (RT) ← (SPR(SPRN)).		355
mr	RT, RS	Move register. (RT) ← (RS) <i>Extended mnemonic for</i> <b>or RT,RS,RS</b>		384
mr.		<i>Extended mnemonic for</i> <b>or. RT,RS,RS</b>	CR[CR0]	
msync		Synchronization. All instructions that precede <b>msync</b> complete before any instructions that follow <b>msync</b> begin. When <b>msync</b> completes, all storage accesses initiated prior to <b>msync</b> will have completed.		358
mtdcr	RS	Move to Condition Register. <i>Extended mnemonic for</i> <b>mtdcr 0xFF,RS</b>		359
mtcrf	FXM, RS	Move some or all of the contents of RS into CR as specified by FXM field, $\text{mask} \leftarrow {}^4(\text{FXM}_0) \parallel {}^4(\text{FXM}_1) \parallel \dots \parallel {}^4(\text{FXM}_6) \parallel {}^4(\text{FXM}_7).$ (CR) ← ((RS) ∧ mask) ∨ (CR) ∧ ¬mask).		359
mtmsr	DCRN, RS	Move to DCR from RS, (DCR(DCRN)) ← (RS).		360
mtmsr	RS	Move to MSR from RS, (MSR) ← (RS).		361

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mtccr0 mtccr1 mtcsrr0 mtcsrr1 mtctr mtdac1 mtdac2 mtdbcr0 mtdbcr1 mtdbcr2 mtdbdr mtdbsr mtdear mtdec mtdecar mtdnv0 mtdnv1 mtdnv2 mtdnv3 mtdtv0 mtdtv1 mtdtv2 mtdtv3 mtdvc1 mtdvc2 mtdvlim mtesr mtiac1 mtiac2 mtiac3 mtiac4 mtinv0 mtinv1 mtinv2 mtinv3 mtitv0 mtitv1 mtitv2 mtitv3 mtivlim mtivor0 mtivor1 mtivor2 mtivor3 mtivor4 mtivor5 mtivor6 mtivor7 mtivor8 mtivor9 mtivor10 mtivor11 mtivor12 mtivor13 mtivor14 mtivor15 mtivpr mtlr mtmcsr mtmcsrr0 mtmcsrr1 mtmmucr mtpid	RS	Move to SPR SPRN. Extended mnemonic for <b>mtspr SPRN,RS</b>  See Table 10-2 Special Purpose Registers Sorted by SPR Number on page 443 for listing of valid SPRN values.		362

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mtsprg0 mtsprg1 mtsprg2 mtsprg3 mtsprg4 mtsprg5 mtsprg6 mtsprg7 mtsrr0 mtsrr1 mttbl mttbu mtcr mtsr mtusprg0 mtixer				
<b>mtspr</b>	SPRN, RS	Move to SPR from RS, (SPR(SPRN)) ← (RS).		362
<b>mulchw</b> <b>mulchw.</b>	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ (signed)	CR[CR0]	365
<b>mulchwu</b> <b>mulchwu.</b>	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ (unsigned)	CR[CR0]	366
<b>mulhww</b> <b>mulhww.</b>	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ (signed)	CR[CR0]	367
<b>mulhwwu</b> <b>mulhwwu.</b>	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ (unsigned)	CR[CR0]	368
<b>mulhw</b> <b>mulhw.</b>	RT, RA, RB	Multiply (RA) and (RB), signed. Place high-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). (RT) ← $prod_{0:31}$ .	CR[CR0]	369
<b>mulhww</b> <b>mulhwwu.</b>	RT, RA, RB	Multiply (RA) and (RB), unsigned. Place high-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (unsigned). (RT) ← $prod_{0:31}$ .	CR[CR0]	370
<b>mullhw</b> <b>mullhw.</b>	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ (signed)	CR[CR0]	371
<b>mullhww</b> <b>mullhwwu.</b>	RT, RA, RB	$(RT)_{16:31} \leftarrow (RA)_{0:15} \times (RB)_{16:31}$ (unsigned)	CR[CR0]	372
<b>mulli</b>	RT, RA, IM	Multiply (RA) and IM, signed. Place low-order result in RT. $prod_{0:47} \leftarrow (RA) \times IM$ (signed) (RT) ← $prod_{16:47}$		373
<b>mullw</b> <b>mullw.</b> <b>mullwo</b> <b>mullwo.</b>	RT, RA, RB	Multiply (RA) and (RB), signed. Place low-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). (RT) ← $prod_{32:63}$ .	CR[CR0] XER[SO, OV] CR[CR0] XER[SO, OV]	374
<b>nand</b> <b>nand.</b>	RA, RS, RB	NAND (RS) with (RB). Place result in RA.	CR[CR0]	375

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>neg</b>	RT, RA	Negative (two's complement) of RA. (RT) $\leftarrow \neg(RA) + 1$		376
<b>neg.</b>			CR[CR0]	
<b>nego</b>			XER[SO, OV]	
<b>nego.</b>			CR[CR0] XER[SO, OV]	
<b>nmacchw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ (RT) $\leftarrow temp_{1:32}$		377
<b>nmacchw.</b>			CR[CR0]	
<b>nmacchwo</b>			XER[SO, OV]	
<b>nmacchwo.</b>			CR[CR0] XER[SO, OV]	
<b>nmacchws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ if (( $prod_0 = RT_0$ ) $\wedge$ ( $RT_0 \neq temp_1$ )) then (RT) $\leftarrow (RT_0 \vee {}^{31}(\neg RT_0))$ else (RT) $\leftarrow temp_{1:32}$		378
<b>nmacchws.</b>			CR[CR0]	
<b>nmacchwso</b>			XER[SO, OV]	
<b>nmacchwso.</b>			CR[CR0] XER[SO, OV]	
<b>nmachhw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ (RT) $\leftarrow temp_{1:32}$		379
<b>nmachhw.</b>			CR[CR0]	
<b>nmachhwo</b>			XER[SO, OV]	
<b>nmachhwo.</b>			CR[CR0] XER[SO, OV]	
<b>nmachhws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ if (( $prod_0 = RT_0$ ) $\wedge$ ( $RT_0 \neq temp_1$ )) then (RT) $\leftarrow (RT_0 \vee {}^{31}(\neg RT_0))$ else (RT) $\leftarrow temp_{1:32}$		380
<b>nmachhws.</b>			CR[CR0]	
<b>nmachhwso</b>			XER[SO, OV]	
<b>nmachhwso.</b>			CR[CR0] XER[SO, OV]	
<b>nmaclhw</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ (RT) $\leftarrow temp_{1:32}$		381
<b>nmaclhw.</b>			CR[CR0]	
<b>nmaclhwo</b>			XER[SO, OV]	
<b>nmaclhwo.</b>			CR[CR0] XER[SO, OV]	
<b>nmaclhws</b>	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ $temp_{0:32} \leftarrow -prod_{0:31} + (RT)$ if (( $prod_0 = RT_0$ ) $\wedge$ ( $RT_0 \neq temp_1$ )) then (RT) $\leftarrow (RT_0 \vee {}^{31}(\neg RT_0))$ else (RT) $\leftarrow temp_{1:32}$		382
<b>nmaclhws.</b>			CR[CR0]	
<b>nmaclhwso</b>			XER[SO, OV]	
<b>nmaclhwso.</b>			CR[CR0] XER[SO, OV]	
<b>nop</b>		Preferred no-op, triggers optimizations based on no-ops. <i>Extended mnemonic for</i> <b>ori 0,0,0</b>		386
<b>nor</b>	RA, RS, RB	NOR (RS) with (RB). Place result in RA.		383
<b>nor.</b>			CR[CR0]	
<b>not</b>	RA, RS	Complement register. (RA) $\leftarrow \neg(RS)$ <i>Extended mnemonic for</i> <b>nor RA,RS,RS</b>		383
<b>not.</b>		<i>Extended mnemonic for</i> <b>nor. RA,RS,RS</b>	CR[CR0]	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>or</b>	RA, RS, RB	OR (RS) with (RB).	CR[CR0]	384
<b>or.</b>		Place result in RA.		
<b>orc</b>	RA, RS, RB	OR (RS) with $\neg$ (RB).	CR[CR0]	385
<b>orc.</b>		Place result in RA.		
<b>ori</b>	RA, RS, IM	OR (RS) with ( $^{16}0 \parallel$ IM). Place result in RA.		386
<b>oris</b>	RA, RS, IM	OR (RS) with (IM $\parallel$ $^{16}0$ ). Place result in RA.		387
<b>rfdi</b>		Return from critical interrupt (PC) $\leftarrow$ (CSRR0). (MSR) $\leftarrow$ (CSRR1).		388
<b>rfdi</b>		Return from interrupt. (PC) $\leftarrow$ (SRR0). (MSR) $\leftarrow$ (SRR1).		389
<b>rfmci</b>		Return from machine check interrupt (PC) $\leftarrow$ (MCSRR0). (MSR) $\leftarrow$ (MCSRR1).		390
<b>rlwimi</b>	RA, RS, SH, MB, ME	Rotate left word immediate, then insert according to mask. $r \leftarrow$ ROTL((RS), SH) $m \leftarrow$ MASK(MB, ME) (RA) $\leftarrow$ ( $r \wedge m$ ) $\vee$ ((RA) $\wedge$ $\neg m$ )	CR[CR0]	391
<b>rlwimi.</b>				
<b>rlwinm</b>	RA, RS, SH, MB, ME	Rotate left word immediate, then AND with mask. $r \leftarrow$ ROTL((RS), SH) $m \leftarrow$ MASK(MB, ME) (RA) $\leftarrow$ ( $r \wedge m$ )	CR[CR0]	392
<b>rlwinm.</b>				
<b>rlwnm</b>	RA, RS, RB, MB, ME	Rotate left word, then AND with mask. $r \leftarrow$ ROTL((RS), (RB) <sub>27:31</sub> ) $m \leftarrow$ MASK(MB, ME) (RA) $\leftarrow$ ( $r \wedge m$ )	CR[CR0]	394
<b>rlwnm.</b>				
<b>rotlw</b>	RA, RS, RB	Rotate left. (RA) $\leftarrow$ ROTL((RS), (RB) <sub>27:31</sub> ) <i>Extended mnemonic for</i> <b>rlwnm RA,RS,RB,0,31</b>	CR[CR0]	394
<b>rotlw.</b>		<i>Extended mnemonic for</i> <b>rlwnm. RA,RS,RB,0,31</b>		
<b>rotlwi</b>	RA, RS, n	Rotate left immediate. (RA) $\leftarrow$ ROTL((RS), n) <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,0,31</b>	CR[CR0]	392
<b>rotlwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,0,31</b>		
<b>rotrwi</b>	RA, RS, n	Rotate right immediate. (RA) $\leftarrow$ ROTL((RS), 32-n) <i>Extended mnemonic for</i> <b>rlwinm RA,RS,32-n,0,31</b>	CR[CR0]	392
<b>rotrwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,32-n,0,31</b>		
<b>sc</b>		System call exception is generated. (SRR1) $\leftarrow$ (MSR) (SRR0) $\leftarrow$ (PC) PC $\leftarrow$ EVPR <sub>0:15</sub> $\parallel$ 0x0C00 (MSR[WE, PR, EE, PE, DR, IR]) $\leftarrow$ 0		395

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>slw</b>	RA, RS, RB	Shift left (RS) by (RB) <sub>27:31</sub> . $n \leftarrow (RB)_{27:31}$ . $r \leftarrow \text{ROTL}((RS), n)$ . if (RB) <sub>26</sub> = 0 then $m \leftarrow \text{MASK}(0, 31 - n)$ else $m \leftarrow 32_0$ (RA) $\leftarrow r \wedge m$ .	CR[CR0]	396
<b>slw.</b>				
<b>slwi</b>	RA, RS, n	Shift left immediate. ( $n < 32$ ) (RA) <sub>0:31-n</sub> $\leftarrow$ (RS) <sub>n:31</sub> (RA) <sub>32-n:31</sub> $\leftarrow$ $n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,n,0,31-n</b>	CR[CR0]	392
<b>slwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,n,0,31-n</b>		
<b>sraw</b>	RA, RS, RB	Shift right algebraic (RS) by (RB) <sub>27:31</sub> . $n \leftarrow (RB)_{27:31}$ . $r \leftarrow \text{ROTL}((RS), 32 - n)$ . if (RB) <sub>26</sub> = 0 then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow 32_0$ $s \leftarrow (RS)_0$ (RA) $\leftarrow (r \wedge m) \vee (32_s \wedge \neg m)$ . XER[CA] $\leftarrow s \wedge ((r \wedge \neg m) \neq 0)$ .	CR[CR0]	397
<b>sraw.</b>				
<b>srawi</b>	RA, RS, SH	Shift right algebraic (RS) by SH. $n \leftarrow \text{SH}$ . $r \leftarrow \text{ROTL}((RS), 32 - n)$ . $m \leftarrow \text{MASK}(n, 31)$ . $s \leftarrow (RS)_0$ (RA) $\leftarrow (r \wedge m) \vee (32_s \wedge \neg m)$ . XER[CA] $\leftarrow s \wedge ((r \wedge \neg m) \neq 0)$ .	CR[CR0]	398
<b>srawi.</b>				
<b>srw</b>	RA, RS, RB	Shift right (RS) by (RB) <sub>27:31</sub> . $n \leftarrow (RB)_{27:31}$ . $r \leftarrow \text{ROTL}((RS), 32 - n)$ . if (RB) <sub>26</sub> = 0 then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow 32_0$ (RA) $\leftarrow r \wedge m$ .	CR[CR0]	399
<b>srw.</b>				
<b>srwi</b>	RA, RS, n	Shift right immediate. ( $n < 32$ ) (RA) <sub>n:31</sub> $\leftarrow$ (RS) <sub>0:31-n</sub> (RA) <sub>0:n-1</sub> $\leftarrow$ $n_0$ <i>Extended mnemonic for</i> <b>rlwinm RA,RS,32-n,n,31</b>	CR[CR0]	392
<b>srwi.</b>		<i>Extended mnemonic for</i> <b>rlwinm. RA,RS,32-n,n,31</b>		
<b>stb</b>	RS, D(RA)	Store byte (RS) <sub>24:31</sub> in memory at EA = (RA 0) + EXTS(D).		400
<b>stbu</b>	RS, D(RA)	Store byte (RS) <sub>24:31</sub> in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) $\leftarrow$ EA.		401
<b>stbux</b>	RS, RA, RB	Store byte (RS) <sub>24:31</sub> in memory at EA = (RA 0) + (RB). Update the base address, (RA) $\leftarrow$ EA.		402
<b>stbx</b>	RS, RA, RB	Store byte (RS) <sub>24:31</sub> in memory at EA = (RA 0) + (RB).		403
<b>sth</b>	RS, D(RA)	Store halfword (RS) <sub>16:31</sub> in memory at EA = (RA 0) + EXTS(D).		404



Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>sthbrx</b>	RS, RA, RB	Store halfword (RS) <sub>16:31</sub> byte-reversed in memory at EA = (RA 0) + (RB). MS(EA, 2) ← (RS) <sub>24:31</sub>    (RS) <sub>16:23</sub>		405
<b>sthu</b>	RS, D(RA)	Store halfword (RS) <sub>16:31</sub> in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) ← EA.		406
<b>sthux</b>	RS, RA, RB	Store halfword (RS) <sub>16:31</sub> in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		407
<b>sthx</b>	RS, RA, RB	Store halfword (RS) <sub>16:31</sub> in memory at EA = (RA 0) + (RB).		408
<b>stmw</b>	RS, D(RA)	Store consecutive words from RS through GPR(31) in memory starting at EA = (RA 0) + EXTS(D).		409
<b>stswi</b>	RS, RA, NB	Store consecutive bytes in memory starting at EA=(RA 0). Number of bytes n=32 if NB=0, else n=NB. Bytes are unstacked from CEIL(n/4) consecutive registers starting with RS. GPR(0) is consecutive to GPR(31).		409
<b>stswx</b>	RS, RA, RB	Store consecutive bytes in memory starting at EA=(RA 0)+(RB). Number of bytes n=XER[TBC]. Bytes are unstacked from CEIL(n/4) consecutive registers starting with RS. GPR(0) is consecutive to GPR(31).		411
<b>stw</b>	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXTS(D).		412
<b>stwbrx</b>	RS, RA, RB	Store word (RS) byte-reversed in memory at EA = (RA 0) + (RB). MS(EA, 4) ← (RS) <sub>24:31</sub>    (RS) <sub>16:23</sub>    (RS) <sub>8:15</sub>    (RS) <sub>0:7</sub>		413
<b>stwcx.</b>	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB) only if reservation bit is set. if RESERVE = 1 then MS(EA, 4) ← (RS) RESERVE ← 0 (CR[CR0]) ← <sup>2</sup> 0    1    XER <sub>so</sub> else (CR[CR0]) ← <sup>2</sup> 0    0    XER <sub>so</sub> .		414
<b>stwu</b>	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) ← EA.		416
<b>stwux</b>	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		417
<b>stwx</b>	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB).		418

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>sub</b>	RT, RA, RB	Subtract (RB) from (RA). $(RT) \leftarrow \neg(RB) + (RA) + 1.$ <i>Extended mnemonic for</i> <b>subf RT,RB,RA</b>		419
<b>sub.</b>		<i>Extended mnemonic for</i> <b>subf. RT,RB,RA</b>	CR[CR0]	
<b>subo</b>		<i>Extended mnemonic for</i> <b>subfo RT,RB,RA</b>	XER[SO, OV]	
<b>subo.</b>		<i>Extended mnemonic for</i> <b>subfo. RT,RB,RA</b>	CR[CR0] XER[SO, OV]	
<b>subc</b>	RT, RA, RB	Subtract (RB) from (RA). $(RT) \leftarrow \neg(RB) + (RA) + 1.$ Place carry-out in XER[CA]. <i>Extended mnemonic for</i> <b>subfc RT,RB,RA</b>		420
<b>subc.</b>		<i>Extended mnemonic for</i> <b>subfc. RT,RB,RA</b>	CR[CR0]	
<b>subco</b>		<i>Extended mnemonic for</i> <b>subfco RT,RB,RA</b>	XER[SO, OV]	
<b>subco.</b>		<i>Extended mnemonic for</i> <b>subfco. RT,RB,RA</b>	CR[CR0] XER[SO, OV]	
<b>subf</b>	RT, RA, RB	Subtract (RA) from (RB). $(RT) \leftarrow \neg(RA) + (RB) + 1.$		419
<b>subf.</b>			CR[CR0]	
<b>subfo</b>			XER[SO, OV]	
<b>subfo.</b>			CR[CR0] XER[SO, OV]	
<b>subfc</b>	RT, RA, RB	Subtract (RA) from (RB). $(RT) \leftarrow \neg(RA) + (RB) + 1.$ Place carry-out in XER[CA].		420
<b>subfc.</b>			CR[CR0]	
<b>subfco</b>			XER[SO, OV]	
<b>subfco.</b>			CR[CR0] XER[SO, OV]	
<b>subfe</b>	RT, RA, RB	Subtract (RA) from (RB) with carry-in. $(RT) \leftarrow \neg(RA) + (RB) + XER[CA].$ Place carry-out in XER[CA].		421
<b>subfe.</b>			CR[CR0]	
<b>subfeo</b>			XER[SO, OV]	
<b>subfeo.</b>			CR[CR0] XER[SO, OV]	
<b>subfic</b>	RT, RA, IM	Subtract (RA) from EXTS(IM). $(RT) \leftarrow \neg(RA) + EXTS(IM) + 1.$ Place carry-out in XER[CA].		422
<b>subfme</b>	RT, RA, RB	Subtract (RA) from (-1) with carry-in. $(RT) \leftarrow \neg(RA) + (-1) + XER[CA].$ Place carry-out in XER[CA].		423
<b>subfme.</b>			CR[CR0]	
<b>subfmeo</b>			XER[SO, OV]	
<b>subfmeo.</b>			CR[CR0] XER[SO, OV]	
<b>subfze</b>	RT, RA, RB	Subtract (RA) from zero with carry-in. $(RT) \leftarrow \neg(RA) + XER[CA].$ Place carry-out in XER[CA].		424
<b>subfze.</b>			CR[CR0]	
<b>subfzeo</b>			XER[SO, OV]	
<b>subfzeo.</b>			CR[CR0] XER[SO, OV]	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>subi</b>	RT, RA, IM	Subtract EXTS(IM) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> <b>addi RT,RA,-IM</b>		251
<b>subic</b>	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> <b>addic RT,RA,-IM</b>		252
<b>subic.</b>	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> <b>addic. RT,RA,-IM</b>	CR[CR0]	253
<b>subis</b>	RT, RA, IM	Subtract (IM    <sup>16</sup> 0) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> <b>addis RT,RA,-IM</b>		254
<b>tlbre</b>	RT, RA, WS	$tlbentry \leftarrow TLB[(RA)_{26:31}]$ if WS = 0 $(RT)_{0:27} \leftarrow tlbentry[EPN,V,TS,SIZE]$ $(RT)_{28:31} \leftarrow 40$ $MMUCR[STID] \leftarrow tlbentry[TID]$ else if WS = 1 $(RT)_{0:21} \leftarrow tlbentry[RPN]$ $(RT)_{22:27} \leftarrow 00$ $(RT)_{28:31} \leftarrow tlbentry[ERP]$ else if WS = 2 $(RT)_{0:15} \leftarrow 160$ $(RT)_{16:24} \leftarrow tlbentry[U0,U1,U2,U3,W,I,M,G,E]$ $(RT)_{25} \leftarrow 0$ $(RT)_{26:31} \leftarrow tlbentry[UX,UW,UR,SX,SW,SR]$ else (RT), MMUCR[STID] $\leftarrow$ undefined		425
<b>tlbsx</b>	RT,RA,RB	Search the TLB for a valid entry that translates the EA. EA = (RA 0) + (RB) if Rc = 1 CR[CR0] <sub>0</sub> $\leftarrow$ 0 CR[CR0] <sub>1</sub> $\leftarrow$ 0 CR[CR0] <sub>3</sub> $\leftarrow$ XER[SO] if Valid TLB entry matching EA and MMUCR[STID,STS] is in the TLB then (RT) $\leftarrow$ Index of matching TLB Entry if Rc = 1 CR[CR0] <sub>2</sub> $\leftarrow$ 1 else (RT) $\leftarrow$ Undefined if Rc = 1 CR[CR0] <sub>2</sub> $\leftarrow$ 0	CR[CR0]	427
<b>tlbsx.</b>				
<b>tlbsync</b>		<b>tlbsync</b> does not complete until all previous TLB-update instructions executed by this processor have been received and completed by all other processors. For the PPC440x5 core, <b>tlbsync</b> is a no-op.		428

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>tlbwe</b>	RS, RA, WS	tlbentry $\leftarrow$ TLB[(RA) <sub>26:31</sub> ] if WS = 0 tlbentry[EPN, V, TS, SIZE] $\leftarrow$ (RS) <sub>0:27</sub> tlbentry[TID] $\leftarrow$ MMUCR[STID] else if WS = 1 tlbentry[RPN] $\leftarrow$ (RS) <sub>0:21</sub> tlbentry[ERP] $\leftarrow$ (RS) <sub>28:31</sub> else if WS = 2 tlbentry[U0, U1, U2, U3, W, I, M, G, E] $\leftarrow$ (RS) <sub>16:24</sub> tlbentry[UX, UW, UR, SX, SW, SR] $\leftarrow$ (RS) <sub>26:31</sub> else tlbentry $\leftarrow$ undefined		429

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>trap</b>		Trap unconditionally. <i>Extended mnemonic for</i> <b>tw 31,0,0</b>		430
<b>tweq</b>	RA, RB	Trap if (RA) equal to (RB). <i>Extended mnemonic for</i> <b>tw 4,RA,RB</b>		
<b>twge</b>		Trap if (RA) greater than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 12,RA,RB</b>		
<b>twgt</b>		Trap if (RA) greater than (RB). <i>Extended mnemonic for</i> <b>tw 8,RA,RB</b>		
<b>twle</b>		Trap if (RA) less than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 20,RA,RB</b>		
<b>twlge</b>		Trap if (RA) logically greater than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 5,RA,RB</b>		
<b>twlgt</b>		Trap if (RA) logically greater than (RB). <i>Extended mnemonic for</i> <b>tw 1,RA,RB</b>		
<b>twlle</b>		Trap if (RA) logically less than or equal to (RB). <i>Extended mnemonic for</i> <b>tw 6,RA,RB</b>		
<b>twllt</b>		Trap if (RA) logically less than (RB). <i>Extended mnemonic for</i> <b>tw 2,RA,RB</b>		
<b>twlng</b>		Trap if (RA) logically not greater than (RB). <i>Extended mnemonic for</i> <b>tw 6,RA,RB</b>		
<b>twlnl</b>		Trap if (RA) logically not less than (RB). <i>Extended mnemonic for</i> <b>tw 5,RA,RB</b>		
<b>twlt</b>		Trap if (RA) less than (RB). <i>Extended mnemonic for</i> <b>tw 16,RA,RB</b>		
<b>twne</b>		Trap if (RA) not equal to (RB). <i>Extended mnemonic for</i> <b>tw 24,RA,RB</b>		
<b>twng</b>		Trap if (RA) not greater than (RB). <i>Extended mnemonic for</i> <b>tw 20,RA,RB</b>		
<b>twnl</b>		Trap if (RA) not less than (RB). <i>Extended mnemonic for</i> <b>tw 12,RA,RB</b>		
<b>tw</b>		TO, RA, RB	Trap exception is generated if, comparing (RA) with (RB), any condition specified by TO is true.	

Table A-1. PPC440x5 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
<b>tweqi</b>	RA, IM	Trap if (RA) equal to EXTS(IM). <i>Extended mnemonic for</i> <b>wi 4,RA,IM</b>		433
<b>twgei</b>		Trap if (RA) greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 12,RA,IM</b>		
<b>twgti</b>		Trap if (RA) greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 8,RA,IM</b>		
<b>twlei</b>		Trap if (RA) less than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 20,RA,IM</b>		
<b>twlgei</b>		Trap if (RA) logically greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>wi 5,RA,IM</b>		
<b>twlgti</b>		Trap if (RA) logically greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 1,RA,IM</b>		
<b>twllei</b>		Trap if (RA) logically less than or equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 6,RA,IM</b>		
<b>twllti</b>		Trap if (RA) logically less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 2,RA,IM</b>		
<b>twlngi</b>		Trap if (RA) logically not greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 6,RA,IM</b>		
<b>twlnli</b>		Trap if (RA) logically not less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 5,RA,IM</b>		
<b>twlti</b>		Trap if (RA) less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 16,RA,IM</b>		
<b>twnei</b>		Trap if (RA) not equal to EXTS(IM). <i>Extended mnemonic for</i> <b>twi 24,RA,IM</b>		
<b>twngi</b>		Trap if (RA) not greater than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 20,RA,IM</b>		
<b>twnli</b>		Trap if (RA) not less than EXTS(IM). <i>Extended mnemonic for</i> <b>twi 12,RA,IM</b>		
<b>twi</b>	TO, RA, IM	Trap exception is generated if, comparing (RA) with EXTS(IM), any condition specified by TO is true.		433
<b>wrtee</b>	RS	Write value of RS <sub>16</sub> to MSR[EE].		435
<b>wrteei</b>	E	Write value of E to MSR[EE].		436
<b>xor</b> <b>xor.</b>	RA, RS, RB	XOR (RS) with (RB). Place result in RA.	CR[CR0]	437
<b>xori</b>	RA, RS, IM	XOR (RS) with ( <sup>16</sup> 0    IM). Place result in RA.		438
<b>xoris</b>	RA, RS, IM	XOR (RS) with (IM    <sup>16</sup> 0). Place result in RA.		439

### A.3 Allocated Instruction Opcodes

Allocated instructions are provided for purposes that are outside the scope of PowerPC Book-E architecture, and are for implementation-dependent and application-specific use, including use within auxiliary processors.

Table A-2 lists the blocks of opcodes which have been allocated by PowerPC Book-E for these purposes. In the table, the character “u” designates a secondary opcode bit which can be set to any value. In some cases, the decimal value of a secondary opcode is shown in parentheses after the binary value.

Table A-2. Allocated Opcodes

Primary Opcode	Extended Opcodes	PPC440x5 Usage
0	All instruction encodings (bits 6:31) except 0x00000000 (the instruction encoding of 0x00000000 is and always will be reserved-illegal)	None
4	All instruction encodings (bits 6:31)	Various (see Table A-5 on page 545)
19	Secondary opcodes (bits 21:30) = 0b000000u11u	None
31	Secondary opcodes (bits 21:30) = 0b000000011u Secondary opcodes (bits 21:30) = 0b000000u110 Secondary opcode (bits 21:30) = 0b0101010110 (342) Secondary opcode (bits 21:30) = 0b0101110110 (374) Secondary opcode (bits 21:30) = 0b1100110110 (822)	Various (see Table A-5 on page 545)
59	Secondary opcodes (bits 21:30) = 0b000000u10u	None
63	Secondary opcodes (bits 21:30) = 0b000000u10u (except secondary opcode decimal 12, which is the <b>fsrp</b> defined instruction)	None

All of the allocated opcodes listed in the table above are available for use by auxiliary processors attached to the PPC440x5, except for those which have already been implemented within the PPC440x5 core for certain implementation-specific purposes. As indicated in the table above, this is the case for certain secondary opcodes within primary opcodes 4 and 31. These opcodes are identified in Table A-5 on page 545, along with all of the defined, preserved, and reserved-nop class opcodes which are implemented within the PPC440x5 core.

### A.4 Preserved Instruction Opcodes

The preserved instruction class is provided to support backward compatibility with the PowerPC Architecture, and/or earlier versions of the PowerPC Book-E architecture. This instruction class includes opcodes which were defined for these previous architectures, but which are no longer defined for PowerPC Book-E.

Table A-3 lists the reserved opcodes designated by PowerPC Book-E. The decimal value of the secondary opcode is shown in parentheses after the binary value.

Table A-3. Preserved Opcodes

Primary Opcode	Extended Opcode	Preserved Mnemonic	PPC440x5 Usage
31	0b0011010010 (210)	mtsr	
31	0b0011110010 (242)	mtsrin	
31	0b0101110010 (370)	tlbia	
31	0b0100110010 (306)	tlbie	
31	0b0101110011 (371)	mftb	Yes

Table A-3. Preserved Opcodes (continued)

Primary Opcode	Extended Opcode	Preserved Mnemonic	PPC440x5 Usage
31	0b1001010011 (595)	mfsr	
31	0b1010010011 (659)	mfsrin	
31	0b0100110110 (310)	eciwx	
31	0b0110110110 (438)	ecowx	

As indicated in the table above, the only preserved opcode which is implemented within the PPC440x5 core is the **mftb** instruction. See *Preserved Instruction Class* on page 54 for more information on PPC440x5 support for this instruction. All other preserved instructions are treated as reserved by PPC440x5 and will cause Illegal Instruction exception type Program interrupts if their execution is attempted.

The preserved opcode for **mftb** is included in *Table A-5* on page 545, along with all of the defined, allocated, and reserved-nop class opcodes which are implemented within the PPC440x5 core.

## A.5 Reserved Instruction Opcodes

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

Reserved instructions are available for future versions of PowerPC Book-E architecture. That is, future versions of PowerPC 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.

Table A-4 lists the reserved-nop opcodes designated by PowerPC Book-E. In the table, the character “u” designates a secondary opcode bit which can be set to any value. All other reserved opcodes are in the reserved-illegal class.

Table A-4. Reserved-nop Opcodes

Primary Opcode	Extended Opcode
31	0b10uuu10010

As shown in the table, there are a total of eight (8) secondary opcodes in the reserved-nop class. The PPC440x5 implements all of the reserved-nop instruction opcodes as true no-ops. These opcodes are included in *Table A-5* on page 545, along with all of the defined, allocated, and preserved class opcodes which are implemented within the PPC440x5 core.

## A.6 Implemented Instructions Sorted by Opcode

*Table A-5* on page 545 lists all of the instructions which have been implemented within the PPC440x5 core, sorted by primary and secondary opcode. These include defined, allocated, preserved, and reserved-nop class instructions (see *Instruction Classes* on page 52 for a more detailed description of each of these instruction classes). Opcodes which are *not* implemented in the PPC440x5 core are *not* shown in the table, and consist of the following:

- Defined instructions



These include the floating-point operations (which may be implemented in an auxiliary processor and executed via the AP interface), as well as the 64-bit operations and the **tlbiva** and **mfapidi** instructions, all of which are handled as reserved-illegal instructions by the PPC440x5.

- Allocated instructions

These include all of the allocated opcodes identified in *Table A-2* on page 543 which are not already implemented within the PPC440x5 core. If not implemented within an attached auxiliary processor, these instructions will be handled as reserved-illegal by the PPC440x5.

- Preserved instructions

These include all of the preserved opcodes identified in *Table A-3* on page 543 except for the **mftb** opcode (which *is* implemented and thus included in *Table A-5*). These instructions will be handled as reserved-illegal by the PPC440x5.

- Reserved instructions

These include all of the reserved opcodes as defined by *Appendix A.5* on page 544, except for the reserved-nop opcodes identified in *Table A-4* on page 544. These instructions by definition are all in the reserved-illegal class and will be handled as such by the PPC440x5.

All PowerPC Book-E instructions are four bytes long and word aligned. All instructions have a primary opcode field (shown as field OPCD in *Figure A-1* through *Figure A-9*, beginning on page 510) in bits 0:5. Some instructions also have a secondary opcode field (shown as field XO in *Figure A-1* through *Figure A-9*).

The “Form” indicated in the table refers to the arrangement of valid field combinations within the four-byte instruction. See *Appendix A.1* on page 507, for the field layouts of each form.

Form X has a 10-bit secondary opcode field, while form XO uses only the low-order 9-bits of that field. Form XO uses the high-order secondary opcode bit (the tenth bit) as a variable; therefore, every form XO instruction really consumes two secondary opcodes from the 10-bit secondary-opcode space. The implicitly consumed secondary opcode is listed in parentheses for form XO instructions in the table below.

*Table A-5. PPC440x5 Instructions by Opcode*

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
3		D	<b>twi</b>	TO, RA, IM	433
4	8	X	mulhhuw mulhhuw.	RT, RA, RB	368
4	12 (524)	XO	machhuw machhuw. machhuwo machhuwo.	RT, RA, RB	343
4	40	X	mulhhw mulhhw.	RT, RA, RB	367
4	44 (556)	XO	machhw machhw. machhwo machhwo.	RT, RA, RB	340

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
4	46 (558)	XO	nmachhw	RT, RA, RB	379
			nmachhw.		
			nmachhwo		
			nmachhwo.		
4	76 (588)	XO	machwsu	RT, RA, RB	342
			machwsu.		
			machwsuo		
			machwsuo.		
4	108 (620)	XO	machhws	RT, RA, RB	341
			machhws.		
			machhwso		
			machhwso.		
4	110 (622)	XO	nmachhws	RT, RA, RB	380
			nmachhws.		
			nmachhwso		
			nmachhwso.		
4	136	X	mulchwu	RT, RA, RB	366
			mulchwu.		
4	140 (652)	XO	macchwu	RT, RA, RB	339
			macchwu.		
			macchwuo		
			macchwuo.		
4	168	X	mulchw	RT, RA, RB	365
			mulchw.		
4	172 (684)	XO	macchw	RT, RA, RB	336
			macchw.		
			macchwo		
			macchwo.		
4	174 (686)	XO	nmacchw	RT, RA, RB	377
			nmacchw.		
			nmacchwo		
			nmacchwo.		
4	204 (716)	XO	macchwsu	RT, RA, RB	338
			macchwsu.		
			macchwsuo		
			macchwsuo.		
4	236 (748)	XO	macchws	RT, RA, RB	337
			macchws.		
			macchwso		
			macchwso.		
4	238 (750)	XO	nmacchws	RT, RA, RB	378
			nmacchws.		
			nmacchwso		
			nmacchwso.		

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
4	392	X	mullhwu	RT, RA, RB	372
			mullhwu.		
4	396 (908)	XO	machwu	RT, RA, RB	347
			machwu.		
			machwuo		
			machwuo.		
4	424	X	mullhw	RT, RA, RB	371
			mullhw.		
4	428 (940)	XO	machw	RT, RA, RB	344
			machw.		
			machwo		
			machwo.		
4	430 (942)	XO	nmachw	RT, RA, RB	381
			nmachw.		
			nmachwo		
			nmachwo.		
4	460 (972)	XO	machwsu	RT, RA, RB	346
			machwsu.		
			machwsuo		
			machwsuo.		
4	492 (1004)	XO	machws	RT, RA, RB	345
			machws.		
			machwso		
			machwso.		
4	494 (1006)	XO	nmachws	RT, RA, RB	382
			nmachws.		
			nmachwso		
			nmachwso.		
7		D	<b>mulli</b>	RT, RA, IM	373
8		D	<b>subfic</b>	RT, RA, IM	422
10		D	<b>cmpli</b>	BF, 0, RA, IM	278
11		D	<b>cmpi</b>	BF, 0, RA, IM	276
12		D	<b>addic</b>	RT, RA, IM	252
13		D	<b>addic.</b>	RT, RA, IM	253
14		D	<b>addi</b>	RT, RA, IM	251
15		D	<b>addis</b>	RT, RA, IM	254
16		B	<b>bc</b>	BO, BI, target	262
			<b>bca</b>		
			<b>bcl</b>		
			<b>bcla</b>		
17		SC	<b>sc</b>		395
18		I	<b>b</b>	target	261
			<b>ba</b>		
			<b>bl</b>		
			<b>bla</b>		

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
19	0	XL	<b>mcrf</b>	BF, BFA	350
19	16	XL	<b>bclr</b>	BO, BI	271
			<b>bclrl</b>		
19	33	XL	<b>crnor</b>	BT, BA, BB	284
19	38	XL	<b>rfmci</b>		390
19	50	XL	<b>rfi</b>		389
19	51	XL	<b>rfci</b>		388
19	129	XL	<b>crandc</b>	BT, BA, BB	281
19	150	XL	<b>isync</b>		311
19	193	XL	<b>crxor</b>	BT, BA, BB	287
19	225	XL	<b>crnand</b>	BT, BA, BB	283
19	257	XL	<b>crand</b>	BT, BA, BB	280
19	289	XL	<b>creqv</b>	BT, BA, BB	282
19	417	XL	<b>crorc</b>	BT, BA, BB	286
19	449	XL	<b>cror</b>	BT, BA, BB	285
19	528	XL	<b>bcctr</b>	BO, BI	268
			<b>bcctrl</b>		
20		M	<b>rlwimi</b>	RA, RS, SH, MB, ME	391
			<b>rlwimi.</b>		
21		M	<b>rlwinm</b>	RA, RS, SH, MB, ME	392
			<b>rlwinm.</b>		
23		M	<b>rlwnm</b>	RA, RS, RB, MB, ME	394
			<b>rlwnm.</b>		
24		D	<b>ori</b>	RA, RS, IM	386
25		D	<b>oris</b>	RA, RS, IM	387
26		D	<b>xori</b>	RA, RS, IM	438
27		D	<b>xoris</b>	RA, RS, IM	439
28		D	<b>andi.</b>	RA, RS, IM	259
29		D	<b>andis.</b>	RA, RS, IM	260
31	0	X	<b>cmp</b>	BF, 0, RA, RB	275
31	4	X	<b>tw</b>	TO, RA, RB	430
31	8 (520)	XO	<b>subfc</b>	RT, RA, RB	420
			<b>subfc.</b>		
			<b>subfco</b>		
			<b>subfco.</b>		
31	10 (522)	XO	<b>addc</b>	RT, RA, RB	249
			<b>addc.</b>		
			<b>addco</b>		
			<b>addco.</b>		
31	11 (523)	XO	<b>mulhwu</b>	RT, RA, RB	370
			<b>mulhwu.</b>		
31	15	XO	<b>isel</b>	RT, RA, RB	310
31	19	X	<b>mfcrr</b>	RT	352
31	20	X	<b>lwarx</b>	RT, RA, RB	330
31	22	X	<b>icbt</b>	RA, RB	305

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	23	X	<b>lwzx</b>	RT, RA, RB	335
31	24	X	<b>slw</b>	RA, RS, RB	396
			<b>slw.</b>		
31	26	X	<b>cntlzw</b>	RA, RS	279
			<b>cntlzw.</b>		
31	28	X	<b>and</b>	RA, RS, RB	257
			<b>and.</b>		
31	32	X	<b>cmpl</b>	BF, 0, RA, RB	277
31	40 (552)	XO	<b>subf</b>	RT, RA, RB	419
			<b>subf.</b>		
			<b>subfo</b>		
			<b>subfo.</b>		
31	54	X	<b>dcbst</b>	RA, RB	293
31	55	X	<b>lwzux</b>	RT, RA, RB	334
31	60	X	<b>andc</b>	RA, RS, RB	258
			<b>andc.</b>		
31	75 (587)	XO	<b>mulhw</b>	RT, RA, RB	369
			<b>mulhw.</b>		
31	78	X	<b>dlmzb</b>	RA, RS, RB	300
			<b>dlmzb.</b>		
31	83	X	<b>mfmsr</b>	RT	354
31	86	X	<b>dcbf</b>	RA, RB	289
31	87	X	<b>lbzx</b>	RT, RA, RB	315
31	104 (616)	XO	<b>neg</b>	RT, RA	376
			<b>neg.</b>		
			<b>nego</b>		
			<b>nego.</b>		
31	119	X	<b>lbzux</b>	RT, RA, RB	314
31	124	X	<b>nor</b>	RA, RS, RB	383
			<b>nor.</b>		
31	131	X	<b>wrtee</b>	RS	435
31	136 (648)	XO	<b>subfe</b>	RT, RA, RB	421
			<b>subfe.</b>		
			<b>subfeo</b>		
			<b>subfeo.</b>		
31	138 (650)	XO	<b>adde</b>	RT, RA, RB	250
			<b>adde.</b>		
			<b>addeo</b>		
			<b>addeo.</b>		
31	144	XFX	<b>mtrcf</b>	FXM, RS	359
31	146	X	<b>mtmsr</b>	RS	361
31	150	X	<b>stwcx.</b>	RS, RA, RB	414
31	151	X	<b>stwx</b>	RS, RA, RB	418
31	163	X	<b>wrteei</b>	E	436
31	183	X	<b>stwux</b>	RS, RA, RB	417

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	200 (712)	XO	<b>subfze</b>	RT, RA, RB	424
			<b>subfze.</b>		
			<b>subfzeo</b>		
			<b>subfzeo.</b>		
31	202 (714)	XO	<b>addze</b>	RT, RA	256
			<b>addze.</b>		
			<b>addzeo</b>		
			<b>addzeo.</b>		
31	215	X	<b>stbx</b>	RS, RA, RB	403
31	232 (744)	XO	<b>subfme</b>	RT, RA, RB	423
			<b>subfme.</b>		
			<b>subfmeo</b>		
			<b>subfmeo.</b>		
31	234 (746)	XO	<b>addme</b>	RT, RA	255
			<b>addme.</b>		
			<b>addmeo</b>		
			<b>addmeo.</b>		
31	235 (747)	XO	<b>mullw</b>	RT, RA, RB	374
			<b>mullw.</b>		
			<b>mullwo</b>		
			<b>mullwo.</b>		
31	246	X	<b>dcbtst</b>	RA, RB	293
31	247	X	<b>stbux</b>	RS, RA, RB	401
31	262	X	<b>icbt</b>	RA, RB	305
31	266 (778)	XO	<b>add</b>	RT, RA, RB	248
			<b>add.</b>		
			<b>addo</b>		
			<b>addo.</b>		
31	278	X	<b>dcbt</b>	RA, RB	291
31	279	X	<b>lhzx</b>	RT, RA, RB	324
31	284	X	<b>eqv</b>	RA, RS, RB	301
			<b>eqv.</b>		
31	311	X	<b>lhzux</b>	RT, RA, RB	323
31	316	X	<b>xor</b>	RA, RS, RB	437
			<b>xor.</b>		
31	323	AFX	<b>mfdcr</b>	RT, DCRN	353
31	339	AFX	<b>mfspr</b>	RT, SPRN	355
31	343	X	<b>lhax</b>	RT, RA, RB	319
31	371	AFX	<b>mtfb</b>	RT, SPRN	543
31	375	X	<b>lhaux</b>	RT, RA, RB	318
31	407	X	<b>sthx</b>	RS, RA, RB	408
31	412	X	<b>orc</b>	RA, RS, RB	385
			<b>orc.</b>		
31	439	X	<b>sthux</b>	RS, RA, RB	407

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	444	X	<b>or</b>	RA, RS, RB	384
			<b>or.</b>		
31	451	XFX	<b>mtdcr</b>	DCRN, RS	360
31	454	X	<b>dccci</b>	RA, RB	295
31	459 (971)	XO	<b>divwu</b>	RT, RA, RB	299
			<b>divwu.</b>		
			<b>divwuo</b>		
			<b>divwuo.</b>		
31	467	XFX	<b>mtspr</b>	SPRN, RS	362
31	470	X	<b>dcbi</b>	RA, RB	290
31	476	X	<b>nand</b>	RA, RS, RB	375
			<b>nand.</b>		
31	486	X	<b>dcread</b>	RT, RA, RB	296
31	491 (1003)	XO	<b>divw</b>	RT, RA, RB	298
			<b>divw.</b>		
			<b>divwo</b>		
			<b>divwo.</b>		
31	512	X	<b>mcrxr</b>	BF	351
31	530		Reserved-nop		544
31	533	X	<b>lswx</b>	RT, RA, RB	328
31	534	X	<b>lwbrx</b>	RT, RA, RB	331
31	536	X	<b>srw</b>	RA, RS, RB	399
			<b>srw.</b>		
31	562		Reserved-nop		544
31	566	X	<b>tlbsync</b>		428
31	594		Reserved-nop		544
31	597	X	<b>lswi</b>	RT, RA, NB	326
31	598	X	<b>msync</b>		358
31	626		Reserved-nop		544
31	658		Reserved-nop		544
31	661	X	<b>stswx</b>	RS, RA, RB	411
31	662	X	<b>stwbrx</b>	RS, RA, RB	413
31	690		Reserved-nop		544
31	722		Reserved-nop		544
31	725	X	<b>stswi</b>	RS, RA, NB	409
31	754		Reserved-nop		544
31	758	X	<b>dcba</b>	RA, RB	288
31	790	X	<b>lhbrx</b>	RT, RA, RB	320
31	792	X	<b>sraw</b>	RA, RS, RB	397
			<b>sraw.</b>		
31	824	X	<b>srawi</b>	RA, RS, SH	398
			<b>srawi.</b>		
31	854	X	<b>mbar</b>	MO	348
31	914	X	<b>tlbsx</b>	RT, RA, RB	427
			<b>tlbsx.</b>		

Table A-5. PPC440x5 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	918	X	<b>sthbrx</b>	RS, RA, RB	405
31	922	X	<b>extsh</b>	RA, RS	303
			<b>extsh.</b>		
31	946	X	<b>tlbre</b>	RT, RA, WS	425
31	954	X	<b>extsb</b>	RA, RS	302
			<b>extsb.</b>		
31	966	X	<b>iccci</b>	RA, RB	307
31	978	X	<b>tlbwe</b>	RS, RA, WS	429
31	982	X	<b>icbi</b>	RA, RB	304
31	998	X	<b>icread</b>	RA, RB	308
31	1014	X	<b>dcbz</b>	RA, RB	294
32		D	<b>lwz</b>	RT, D(RA)	332
33		D	<b>lwzu</b>	RT, D(RA)	333
34		D	<b>lbz</b>	RT, D(RA)	312
35		D	<b>lbzu</b>	RT, D(RA)	313
36		D	<b>stw</b>	RS, D(RA)	412
37		D	<b>stwu</b>	RS, D(RA)	416
38		D	<b>stb</b>	RS, D(RA)	400
39		D	<b>stbu</b>	RS, D(RA)	401
40		D	<b>lhz</b>	RT, D(RA)	321
41		D	<b>lhzu</b>	RT, D(RA)	322
42		D	<b>lha</b>	RT, D(RA)	316
43		D	<b>lhau</b>	RT, D(RA)	317
44		D	<b>sth</b>	RS, D(RA)	404
45		D	<b>sthu</b>	RS, D(RA)	406
46		D	<b>lmw</b>	RT, D(RA)	325
47		D	<b>stmw</b>	RS, D(RA)	409



## Appendix B. PPC440x5 Core Compiler Optimizations

This appendix describes some potential optimizations for compilers.

1. Place target addresses (subroutine entry points) on cache line boundaries (32-bytes)
2. Up to five instructions between a load and a use of the load result. Assuming a data cache hit, the worst case scenario for the PPC440x5 core is five instructions between a load-use, in order to avoid any bubbles. The five instructions are:
  - One dispatch, together with the load
  - Two the cycle after
  - Two the cycle after that

In the next cycle, the use of the load result can dispatch. Therefore, the compiler should try to schedule as many as five instructions between the load and use of the load result. However, if some of the instruction pairs between the load-use have pipeline dependencies (such that they cannot dispatch together), there is no benefit in including the extra instructions between the load-use, and other scheduling optimizations could be made.

In the worst case of instruction pairings, the maximum performance can be achieved with only two instructions between the load and use of the load result. This is the case when the load instruction pairs with the instruction before it (instead of after it), and then the next two instructions require the same pipe, so only one can dispatch during the cycle after the load, and then third instruction after the load needs the same pipe as the second, so they cannot dispatch together either. In such a case, the third instruction after the load might as well be the use of the load result. See item 3 for information about which instruction pairings can dispatch together.

3. Pair instructions for dual dispatch. The rules for instruction dispatch in the PPC440x5 core are as follows: loads and stores can only use the L-Pipe. Branches, CR-updates, XER-updates ("o" forms of arithmetic instructions), multiply, divide, system instructions (such as **rfi** and **sc**), and any SPR accesses (**mtspr**, **mfspr**) can only use the I-Pipe. All other instructions (primarily non-CR-updating and non-XER-updating arithmetic and logic instructions) can use either the J-Pipe or the I-Pipe. Instructions should be paired so that they can dispatch as pairs. For example, pair loads and stores with any other instructions. Pair CR-updates with non-CR-updating instructions and so on.
4. Do *not* bother to try to schedule instructions between CR-updates and branches that are conditional on those CR-updates (with some exceptions).

The exceptions are for CR-updates caused by multiply, divide, multiply-accumulate, **mtcrf**, **tbsx.**, and **stwcx.** instructions. If a branch depends on the CR result of one of these instructions, one or more instructions should be scheduled (if possible) between the CR update and the branch. Of course, it is also the general case (as pointed out in item 3) that the compiler should schedule instructions so they can issue in pairs, and a CR-update and a branch both issue to the I-Pipe, so they cannot issue together. (The compiler should try to set things up so a CR-update and a following branch (regardless of any CR-dependency by the branch) can issue in pairs.) This can mean the CR-update can get paired with the instruction before it, and the branch with the instruction after it, such that there is dual issue in both cycles. However, if this pairing is not possible, an instruction should be inserted (if possible, of course; do not create no-ops for no reason) between the CR-update and the branch to allow the dual issue.

The point of this item is to explain that there is no need to separate the CR-update and the branch simply for the sake of the CR-dependency. That is, there is no extra cycle penalty associated with the CR-update/branch CR-dependency, beyond the "standard" penalty of the inability to dual issue, unless the CR-update is one of the types mentioned above.

If the CR-update is MAC or a  $16 \times 32$  multiply, 1 to 3 instructions should be scheduled between the CR-update and the branch (0 or 1 instruction, depending on whether the CR-update pairs with the instruction before or after, or 1 to 2 instructions to issue between the issue of the CR-update and the issue of the branch, depending on whether there is a single-issue or dual-issue opportunity for the instruction(s) which are scheduled between the CR-update and the branch).

Similarly, if the CR-update is  $32 \times 32$  multiply, divide, **tlbsx.**, or **stwcx.**, schedule 3 to 5 instructions between the CR-update and the branch (two issue cycles of 2 to 4 instructions between, plus the 0 to 1 issuing with the CR-update).

Finally, if the CR-update is **mtcrf**, schedule 5 to 7 instructions between (3 cycles of issue between them).

5. Avoid the use of string/multiple instructions (with some exceptions).

The exceptions have to do with cache effects (more cache misses due to more instructions if you use separate loads/stores instead of a string/multiple), and the specialized behavior of a string, where the bytes are inserted into the more-significant portion of the GPR, in preparation for a "string compare" operation to determine which string is "greater" than another. If the string/multiple is for a relatively small number of registers (or the expansion into discrete loads/stores is known to not have an overall detrimental cache impact), and if a string is being used only for a copy operation and the size is known, performance can be improved by using discrete loads/stores. Essentially, due to hazard determination within the processor, string/multiples impose a couple of cycles of extra, "false" penalty on both the front-end and the back-end. On the other hand, if this penalty is amortized over a large number of registers (say 16 or so), the impact of the extra stalls is probably negligible.

6. Insert 10 or so instructions within a **bdnz** loop (loop unrolling).

7. Put 4 to 8 instructions between **mtlr/mtctr** and **blr/bctr**

8. Put 1 to 3 instructions between  $16 \times 32$  multiply and the use of the result.

9. Put 2 to 5 instructions between  $32 \times 32$  multiply and the use of the result.

10. Use the "without allocate" attribute appropriately on block copy operations, such as calls to the library **memcpy** function, or implicit structure copies.

11. Block move operations. If moving a block of memory using a series of load/store operations, perform the load/store operations in the following order: L1-L2-L3-S1-S2-S3, and repeat. Having the second and third loads between the first load and the first store fills the two-cycle load-use penalty.

## Index

### A

add, 248  
 add., 248  
 addc, 249  
 addc., 249  
 addco, 249  
 addco., 249  
 adde, 250  
 adde., 250  
 addeo, 250  
 addeo., 250  
 addi, 251  
 addic, 252  
 addic., 253  
 addis, 254  
 addme, 255  
 addme., 255  
 addmeo, 255  
 addmeo., 255  
 addo, 248  
 addo., 248  
 addressing, 39  
 addressing modes, 41  
   data storage, 41  
   instruction storage, 41  
 addze, 256  
 addze., 256  
 addzeo, 256  
 addzeo., 256  
 alignment  
   load and store, 114  
 Alignment interrupt, 179  
 alignment interrupts, 179  
 allocated instruction summary, 62  
 allocation  
   data cache line on store miss, 116  
 alphabetical summary of implemented instructions, 512  
 and, 257  
 and., 257  
 andc, 258  
 andc., 258  
 andi., 259  
 andis., 260  
 arithmetic compare, 69  
 arrays, shadow TLB, 147  
 asynchronous interrupt class, 153  
 attributes, storage, 140  
 Auxiliary Processor Unavailable interrupt, 184  
 auxiliary processor unavailable interrupt, 184

### B

b, 261  
 ba, 261  
 bc, 262  
 bca, 262  
 bcctr, 268  
 bcctrl, 268  
 bcl, 262  
 bcla, 262  
 bclr, 271  
 bclrl, 271  
 bctr, 268  
 bctrl, 268  
 bdnz, 263  
 bdnza, 263  
 bdnzf, 263  
 bdnzfa, 263  
 bdnzfl, 263  
 bdnzfla, 263  
 bdnzflr, 272  
 bdnzflrl, 272  
 bdnzl, 263  
 bdnzla, 263  
 bdnzlr, 272  
 bdnzlrl, 272  
 bdnzt, 263  
 bdnzta, 263  
 bdnztl, 263  
 bdnztla, 263  
 bdnztlr, 272  
 bdnztlrl, 272  
 bdz, 263  
 bdza, 263  
 bdzf, 264  
 bdzfa, 264  
 bdzfl, 264  
 bdzfla, 264  
 bdzflr, 272  
 bdzflrl, 272  
 bdzl, 263  
 bdzla, 263  
 bdzlr, 272  
 bdzlrl, 272  
 bdzt, 264  
 bdzta, 264  
 bdztl, 264  
 bdztle, 264  
 bdztlr, 272  
 bdztlrl, 272  
 beq, 264  
 beqa, 264  
 beqctr, 269  
 beqctrl, 269  
 beql, 264  
 beqlr, 273  
 beqlrl, 273

bf, 264  
 bfa, 264  
 bfctr, 269  
 bfctrl, 269  
 bfl, 264  
 bfla, 264  
 bflr, 273  
 bfirl, 273  
 bge, 265  
 bgea, 265  
 bgctrl, 269  
 bgel, 265  
 bgela, 265  
 bgelr, 273  
 bgelrl, 273  
 bgrctr, 269  
 bgt, 265  
 bgta, 265  
 bgtctr, 269  
 bgtctrl, 269  
 bgtl, 265  
 bgtla, 265  
 bgtlr, 273  
 bgtlrl, 273  
 BI field on conditional branches, 63  
 big endian  
     defined, 43  
     structure mapping, 44  
 big endian mapping, 43  
 bl, 261  
 bla, 261  
 ble, 265  
 blea, 265  
 blectr, 269  
 blectrl, 269  
 blel, 265  
 blela, 265  
 blelr, 273  
 blelrl, 273  
 blr, 272  
 blrl, 272  
 blt, 265  
 blta, 265  
 bltctr, 269  
 bltctrl, 269  
 bltl, 265  
 bltla, 265  
 bltlr, 273  
 bltlrl, 273  
 bne, 266  
 bnea, 266  
 bnectr, 269  
 bnectrl, 269  
 bnel, 266  
 bnela, 266  
 bnelr, 273  
 bnelrl, 273  
 bng, 266  
 bnga, 266  
 bngctr, 269  
 bngctrl, 269  
 bngl, 266  
 bngla, 266  
 bnglr, 273  
 bnglrl, 273  
 bnl, 266  
 bnla, 266  
 bnlctr, 270  
 bnlctrl, 270  
 bnll, 266  
 bnlla, 266  
 bnllr, 274  
 bnllrl, 274  
 bns, 266  
 bnsa, 266  
 bnsctr, 270  
 bnsctrl, 270  
 bnsl, 266  
 bnsla, 266  
 bnslr, 274  
 bnsrl, 274  
 bnu, 267  
 bnua, 267  
 bnuctr, 270  
 bnuctrl, 270  
 bnul, 267  
 bnula, 267  
 bnulr, 274  
 bnulrl, 274  
 BO field on conditional branches, 63  
 branch instruction summary, 59  
 branch instructions, exception priorities for, 200  
 branch prediction, 64, 513  
 branch processing, 62  
 branch taken (BRT) debug events, 228  
 branching control  
     BI field on conditional branches, 63  
     BO field on conditional branches, 63  
     branch addressing, 62  
     branch prediction, 64  
     registers, 65  
 bso, 267  
 bsoa, 267  
 bsoctr, 270  
 bsoctrl, 270  
 bsol, 267  
 bsola, 267  
 bsolr, 274  
 bsolrl, 274  
 bt, 267  
 bta, 267  
 btctr, 270

btctrl, 270  
 btl, 267  
 btla, 267  
 btlr, 274  
 btlrl, 274  
 bun, 267  
 buna, 267  
 bunctr, 270  
 bunctrl, 270  
 bunl, 267  
 bunla, 267  
 bunlr, 274  
 bunlrl, 274  
 byte ordering, 42
 

- big endian, defined, 43
- instructions, 44, 45
- little endian, defined, 43
- structure mapping
  - big-endian mapping, 43
  - little endian mapping, 44

## C

cache block, defined, 105  
 cache line
 

- See *also* cache block

 cache line locking, 96  
 cache line replacement policy, 94  
 cache locking transient mechanism, 96  
 cache management instructions
 

- summary
  - data cache, 121
  - instruction cache, 105

 caching inhibited, 141  
 CCR0, 74, 105, 107, 122, 448  
 CCR1, 450  
 change status management, 150  
 clrlswi, 392  
 clrlswi., 392  
 clrlwi, 392  
 clrlwi., 392  
 clrrwi, 392  
 clrrwi., 392  
 cmp, 275  
 cmpi, 276  
 cmpl, 277  
 cmpli, 278  
 cmplw, 277  
 cmplwi, 278  
 cmpw, 275  
 cmpwi, 276, 348  
 cntlzw, 279  
 cntlzw., 279  
 code
 

- self-modifying, 103

coherence
 

- data cache, 120

 coherency
 

- instruction cache, 103

 compare
 

- arithmetic, 69
- logical, 69

 Condition Register. *See also* CR  
 context synchronization, 80  
 control
 

- data cache, 121
- instruction cache, 105

 conventions
 

- notational, 25

 CR, 66, 452
 

- defined

 CR updating instructions, 67  
 instructions
 

- integer
  - CR, 68

 crand, 280  
 crandc, 281  
 crclr, 287  
 creqv, 282  
 Critical Input interrupt, 172  
 critical interrupts, 155  
 Critical Save/Restore Register 0, 162, 163  
 Critical Save/Restore Register 1, 162, 163  
 crmove, 285  
 crnand, 283  
 crnor, 284  
 crnot, 284  
 cror, 285  
 crorc, 286  
 crset, 282  
 crxor, 287  
 CSRR0, 162, 163, 453  
 CSRR1, 162, 163, 454  
 CTR, 66, 455

## D

DAC
 

- debug events
  - applied to instructions that result in multiple storage accesses, 224
  - applied to various instruction types, 224
  - fields, 220
  - overview, 220
  - processing, 223
  - registers
    - DAC1–DAC2, 240
- DAC1–DAC2, 240, 456
- Data Address Compare Register (DAC1), 240
- data address compare *See also* DAC, 220

- data addressing modes, 41
- data cache
  - coherency, 120
- data cache array organization and operation, 93
- data cache controller. *See* DCC
- data cache line allocation on store miss, 116
- data read PLB interface requests
  - PLB interface, 118
- data read requests, 118
- data storage addressing modes, 41
- Data Storage interrupt, 175
- data storage interrupts, 175
- Data TLB Error interrupt, 186
- data TLB error interrupts, 186
- data value compare *See also* DVC, 225
- data write PLB interface requests
  - PLB interface, 119
- data write requests, 119
- DBCR0, 233, 237, 457
- DBCR1, 234, 459
- DBCR2, 461
- DBDR, 241
- DBSR, 238
- dcba
  - operation summary, 121
- dcbf, 289
  - operation summary, 121
- dcbi, 290
  - operation summary, 121
- dcbst, 291
  - operation summary, 121
- dcbt
  - formal description, 292
  - functional description, 122
  - operation summary, 121
- dcbt and dcbtst operation, 122
- dcbtst
  - formal description, 293
  - functional description, 122
  - operation summary, 121
- dcbz, 294
  - operation summary, 121
- DCC (data cache controller)
  - control, 121
  - debug, 121
  - features, 112
  - operations, 113
- dccci, 295
  - operation summary, 121
- DCDBTRH, 123, 466
- DCDBTRL, 123, 467
- dcread
  - functional description, 123, 296
  - operation summary, 121
- DCRs
  - defined, 52
- DEAR, 468
- debug
  - debug cache, 121
  - instruction cache, 105
- debug events
  - BRT, 228
  - DAC, 220
  - DAC fields, 220
  - DVC, 225
  - DVC fields, 226
  - IAC, 216, 229
  - IAC fields, 216
  - ICMP, 229
  - IPRT, 230
  - overview, 215
  - RET, 229
  - summary, 231
  - TRAP, 228
  - UDE, 231
- Debug Interrupt, 188
- debug interrupts, 188
- debug modes
  - debug wait, 214
  - external, 214
  - internal, 213
  - overview, 213
  - trace, 215
- debug wait mode, 214
- debugging
  - debug events, 215
  - debug modes, 213
  - development tool support, 213
  - registers
    - DAC1–DAC2, 240
    - DBCR0, 233, 237
    - DBCR1, 234
    - DBDR, 241
    - DBSR, 238
    - DVC1–DVC2, 240
    - IAC1–IAC4, 239
    - overview, 232
  - reset, 232
  - timer freeze, 232
- DEC, 205, 469
- DECAR, 206, 470
- Decrementer Interrupt, 185
- decrementer interrupts, 185
- device control registers, 52
- Device Control Registers. *See also* DCRs
- direct write to memory, 116
- divw, 298
- divw., 298
- divwo, 298
- divwo., 298
- divwu, 299
- divwu., 299



bfl, 264  
 bfla, 264  
 bflr, 273  
 bflrl, 273  
 bge, 265  
 bgea, 265  
 bgectr, 269  
 bgectrl, 269  
 bgel, 265  
 bgela, 265  
 bgelr, 273  
 bgelrl, 273  
 bgt, 265  
 bgta, 265  
 bgtctr, 269  
 bgtctrl, 269  
 bgtl, 265  
 bgtla, 265  
 bgtlr, 273  
 bgtlrl, 273  
 ble, 265  
 blea, 265  
 blectr, 269  
 blectrl, 269  
 blel, 265  
 blela, 265  
 blelr, 273  
 blelrl, 273  
 blr, 272  
 blrl, 272  
 blt, 265  
 blta, 265  
 bltctr, 269  
 bltctrl, 269  
 bltl, 265  
 bltla, 265  
 bltlr, 273  
 bltlrl, 273  
 bne, 266  
 bnea, 266  
 bnectr, 269  
 bnectrl, 269  
 bnel, 266  
 bnela, 266  
 bnelr, 273  
 bnelrl, 273  
 bng, 266  
 bnga, 266  
 bngctr, 269  
 bngctrl, 269  
 bngl, 266  
 bngla, 266  
 bnglr, 273  
 bnglrl, 273  
 bnl, 266  
 bnla, 266  
 bnlctr, 270  
 bnlctrl, 270  
 bnll, 266  
 bnlla, 266  
 bnllr, 274  
 bnllrl, 274  
 bns, 266  
 bnsa, 266  
 bnsctr, 270  
 bnsctrl, 270  
 bnsi, 266  
 bnsia, 266  
 bnslr, 274  
 bnsrl, 274  
 bnu, 267  
 bnua, 267  
 bnuctr, 270  
 bnuctrl, 270  
 bnul, 267  
 bnula, 267  
 bnulr, 274  
 bnulrl, 274  
 bsalr, 274  
 bso, 267  
 bsoa, 267  
 bsoctr, 270  
 bsoctrl, 270  
 bsol, 267  
 bsola, 267  
 bsolrl, 274  
 bt, 267  
 bta, 267  
 btctr, 270  
 btctrl, 270  
 btl, 267  
 btla, 267  
 btlr, 274  
 btlrl, 274  
 bun, 267  
 buna, 267  
 bunctr, 270  
 bunctrl, 270  
 bunl, 267  
 bunla, 267  
 bunlr, 274  
 bunlrl, 274  
 clrslwi, 392  
 clrslwi., 392  
 clrlwi, 392  
 clrlwi., 392  
 clrrwi, 392  
 clrrwi., 392  
 cmplw, 277  
 cmplwi, 278  
 cmpw, 275  
 cmpwi, 276, 348



crclr, 287  
 crmove, 285  
 crnot, 284  
 crset, 282  
 extlwi, 393  
 extlwi., 393  
 extrwi, 393  
 extrwi., 393  
 for addi, 251  
 for addic, 252  
 for addic., 253  
 for addis, 254  
 for bc, bca, bcl, bcla, 263  
 for bcctr, bcctrl, 268  
 for bclr, bclrl, 272  
 for cmp, 275  
 for cmpi, 276  
 for cmpl, 277  
 for cmpli, 278  
 for creqv, 282  
 for crnor, 284  
 for cror, 285  
 for crxor, 287  
 for mbar 0, 348  
 for mfspr, 356, 363  
 for mtrcf, 359  
 for nor, nor., 383  
 for or, or., 384  
 for ori, 386  
 for rlwimi, rlwimi., 391  
 for rlwinm, rlwinm., 392  
 for rlwnm, rlwnm., 394  
 for subf, subf., subfo, subfo., 419  
 for subfc, subfc., subfco, subfco., 420  
 for tw, 431  
 for twi, 434  
 inslwi, 391  
 inslwi., 391  
 insrwi, 391  
 insrwi., 391  
 li, 251  
 lis, 254  
 mr, 384  
 mr., 384  
 mtrc, 359  
 nop, 386  
 not, 383  
 not., 383  
 rotlw, 394  
 rotlw., 394  
 rotlwi, 393  
 rotlwi., 393  
 rotrwi, 393  
 rotrwi., 393  
 slwi, 393  
 slwi., 393  
 srwi, 393  
 srwi., 393  
 sub, 419  
 sub., 419  
 subc, 420  
 subc., 420  
 subco, 420  
 subco., 420  
 subi, 251  
 subic, 252  
 subic., 253  
 subis, 254  
 subo, 419  
 subo., 419  
 trap, 431  
 tweq, 431  
 tweqi, 434  
 twge, 431  
 twgei, 434  
 twgle, 431  
 twgt, 431  
 twgti, 434  
 twle, 431  
 twlei, 434  
 twlgei, 434  
 twlgt, 431  
 twlgti, 434  
 twlle, 431  
 twllei, 434  
 twllt, 431  
 twllti, 434  
 twlng, 431  
 twlngi, 434  
 twlnl, 431  
 twlnli, 434  
 twlt, 431  
 twlti, 434  
 twne, 431  
 twnei, 434  
 twng, 431  
 twngi, 434  
 twnl, 432  
 twnli, 434  
 external debug mode, 214  
 External Input interrupt, 179  
 external input interrupts, 179  
 extlwi, 393  
 extlwi., 393  
 extrwi, 393  
 extrwi., 393  
 extsb, 302  
 extsb., 302

**F**

## features

- DCC, 112
- ICC, 100

## FIT, 206

## fixed interval timer, 206

## fixed interval timer interrupt, 185

## Fixed-Interval Timer interrupt, 185

## floating point interrupt unavailable interrupts, 183

## floating-point load and store instructions, exception priorities for, 196

## Floating-Point Unavailable interrupt, 183

## freezing the timer facilities, 211

**G**

## G storage attribute, 141

General Purpose Registers. *See also* GPRs

## GPR0-GPR31, 477

## GPRs

- defined, 51

## GPRs, illustrated, 69

## guarded, 141

**H, I, J, K**

## I storage attribute, 141

## IAC

## debug events

- fields, 216
- overview, 216, 229
- processing, 219

## registers

- IAC1-IAC4, 239

## IAC1-IAC4, 473, 478

## IAC1-IAC4, 239

## icbi, 304

- operation summary, 105

## icbt

- formal description, 305
- functional description, 108
- operation summary, 105

## ICC (instruction cache controller)

- control, 105
- debug, 105
- features, 100
- operations, 101

## iccci, 307

- operation summary, 105

## ICDBDR, 479

## ICDBTRH, 480

## ICDBTRL, 481

## icread, 308

- functional description, 109

## operation summary, 105

## implemented instruction set summary, 55

## implicit update, 68

## imprecise interrupts, 154

## inslwi, 391

## inslwi., 391

## insrwi, 391

## insrwi., 391

## instruction

- add, 248
- add., 248
- addc, 249
- addc., 249
- addco, 249
- addco., 249
- adde, 250
- adde., 250
- addeo, 250
- addeo., 250
- addi, 251
- addic, 252
- addic., 253
- addis, 254
- addme, 255
- addme., 255
- addmeo, 255
- addmeo., 255
- addo, 248
- addo., 248
- addze, 256
- addze., 256
- addzeo, 256
- addzeo., 256
- and, 257
- and., 257
- andc, 258
- andc., 258
- andi., 259
- andis., 260
- b, 261
- ba, 261
- bc, 262
- bca, 262
- bcctr, 268
- bcctrl, 268
- bcl, 262
- bcla, 262
- bclr, 271
- bclrl, 271
- bl, 261
- bla, 261
- cmp, 275
- cmpi, 276
- cmpl, 277
- cmpli, 278
- cntlzw, 279

cntlzw., 279  
 crand, 280  
 crandc, 281  
 creqv, 282  
 crnand, 283  
 crnor, 284  
 cror, 285  
 crorc, 286  
 crxor, 287  
 dcbf, 289  
 dcbi, 290  
 dcbst, 291  
 dcbt, 292  
 dcbtst, 293  
 dcbz, 294  
 dccci, 295  
 dcread, 296  
 divw, 298  
 divw., 298  
 divwo, 298  
 divwo., 298  
 divwu, 299  
 divwu., 299  
 divwuo, 299  
 divwuo., 299  
 dlmzb, 300  
 dlmzb., 300  
 eqv, 301  
 eqv., 301  
 extsb, 302  
 extsb., 302  
 icbi, 304  
 icbt, 305  
 iccci, 307  
 icread, 308  
 isel, 310  
 isync, 311  
 lbz, 312  
 lbzu, 313  
 lbzx, 315  
 lha, 316  
 lhau, 317  
 lhax, 319  
 lhbrx, 320  
 lhz, 321  
 lhzu, 322  
 lhzux, 323  
 lhzx, 324  
 lmw, 325  
 lswi, 326  
 lswx, 328  
 lwarx, 330  
 lwz, 332  
 lwzu, 333  
 lwzux, 334  
 lwzx, 335  
 macchw, 336  
 macchws, 337  
 macchwsu, 338  
 macchwu, 339  
 machhw, 340  
 machhwsu, 342  
 machhwu, 343  
 maclhw, 344  
 maclhws, 345, 382  
 maclhwu, 347  
 mbar, 348  
 mcrf, 350  
 mcrxr, 351  
 mfcr, 352  
 mfdcr, 353  
 mfmsr, 354  
 mfspr, 355  
 msync, 358  
 mtrf, 359  
 mtdcr, 360  
 mtspr, 362  
 mulchw, 365  
 mulchwu, 366  
 mulhhw, 367  
 mulhhwu, 368  
 mulhwu, 370  
 mulhwu., 370  
 mullhw, 371  
 mullhwu, 372  
 mulli, 373  
 mullw, 374  
 mullw., 374  
 mullwo, 374  
 mullwo., 374  
 nand, 375  
 nand., 375  
 neg, 376  
 neg., 376  
 nego, 376  
 nego., 376  
 nmacchw, 377  
 nmacchws, 378  
 nmachhw, 379  
 nmachhws, 380  
 nmaclhw, 381  
 nmaclhws, 382  
 nor, 383  
 nor., 383  
 or, 384  
 or., 384  
 orc, 385  
 orc., 385  
 ori, 386  
 oris, 387  
 partially executed, 158  
 rfci, 388

- rfi, 389
- rftci, 390
- rlwimi, 391
- rlwimi., 391
- rlwinm, 392
- rlwinm., 392
- rlwnm, 394
- rlwnm., 394
- sc, 395
- slw, 396
- slw., 396
- sraw, 397
- sraw., 397
- srawi, 398
- srawi., 398
- srw, 399
- srw., 399
- stb, 400
- stbu, 401
- stbux, 402
- stbx, 403
- sth, 404
- sthbrx, 405
- sthu, 406
- sthux, 407
- sthx, 408
- stmw, 409
- stswi, 409
- stw, 412
- stwbrx, 413
- stwcx., 414
- stwu, 416
- stwux, 417
- stwx, 418
- subf, 419
- subf., 419
- subfc, 420
- subfc., 420
- subfco, 420
- subfco., 420
- subfe, 421
- subfe., 421
- subfeo, 421
- subfeo., 421
- subfic, 422
- subfme, 423
- subfme., 423
- subfmeo, 423
- subfmeo., 423
- subfo, 419
- subfo., 419
- subfze, 424
- subfze., 424
- subfzeo, 424
- subfzeo., 424
- tlbre, 425
- tlbsx, 427
- tlbsx., 427
- tlbsync, 428
- tlbwe, 429
- tw, 430
- twi, 433
- wrtee, 435
- wrteei, 436
- xor, 437
- xori, 438
- instruction address compare *See also* IAC, 216, 229
- instruction addressing modes, 41
- instruction cache array organization and operation, 93
- instruction cache coherency, 103
- instruction cache controller. *See* ICC
- instruction cache synonyms, 104
- instruction classes, 52
- instruction complete (ICMP) debug events, 229
- instruction fields, 508
- instruction formats, 244, 507
  - diagrams, 509
- instruction forms, 507, 509
  - B-form, 510
  - D-form, 510
  - I-form, 510
  - M-form, 512
  - SC-form, 510
  - X-form, 511
  - AFX-form, 512
  - XL-form, 512
  - XO-form, 512
- instruction set
  - classes, 52
  - summary
    - allocated instructions, 62
    - branch, 59
    - cache management, 61
    - CR logical, 60
    - integer arithmetic, 57
    - integer compare, 58
    - integer logical, 58
    - integer rotate, 58
    - integer shift, 59
    - integer storage access, 56
    - integer trap, 58
    - processor synchronization, 60
    - register management, 60
    - system linkage, 60
    - TLB management, 61
- instruction set portability, 244
- instruction set summary, 55
- instruction storage addressing modes, 41
- Instruction Storage interrupt, 178
- instruction storage interrupts, 178
- Instruction TLB Error Interrupt, 188
- instruction TLB error interrupts, 188

- Instructions
  - classes
    - allocated, 53
  - instructions
    - all other, exception priorities for, 201
    - allocated (other), exception priorities for, 198
    - allocated instruction opcodes, 543
    - allocated load and store, exception priorities for, 197
    - alphabetical listing, 247
    - alphabetical summary, 512
    - branch, exception priorities for, 200
    - byte ordering, 44, 45
    - byte-reverse, 46
    - categories, 243
      - allocated instruction summary, 62
      - branch, 59
      - integer, 56
      - processor control, 59
      - storage control, 61
      - storage synchronization, 61
    - classes
      - defined, 53, 54
      - preserved, 55
    - CR updating, 67
    - DAC debug events applied to
      - cache management, 224
      - instructions that result in multiple storage accesses, 224
      - lswx, stswx, 224
      - special cases, 224
      - stwcx., 224
      - various, 224
    - data cache management instruction summary, 121
    - DVC debug events applied to
      - cache management, 227
      - instructions that result in multiple storage accesses, 227
      - lswx, stswx, 227
      - special cases, 227
      - stwcx., 227
      - various, 227
    - floating-point (other), exception priorities for, 197
    - floating-point load and store, exception priorities for, 196
    - format diagrams, 509
    - formats, 507
    - forms, 507, 509
    - implemented instruction set summary, 55
    - instruction cache management instruction summary, 105
    - integer compare
      - CR update, 69
    - integer load, store, and cache management, exception priorities for, 196
    - mfmsr, 159
    - mtmsr, 159
    - opcodes, 544
    - partially executed, 158
    - preserved instruction opcodes, 543
    - preserved, exception priorities for, 200
    - privileged, 78
    - privileged instructions, exception priorities for, 199
    - pseudocode operator precedence, 247
    - register usage, 247
    - reserved instruction opcodes, 544
    - reserved, exception priorities for, 201
    - reserved-illegal, 544
    - reserved-nop, 544
    - return from interrupt, exception priorities for, 200
    - rfi, 161
    - sorted by opcode, 544
    - syntax summary, 513
    - system call, exception priorities for, 200
    - trap, exception priorities for, 199
  - integer instructions
    - arithmetic, 57
    - compare, 58
    - logical, 58
    - rotate, 58
    - shift, 59
    - storage access, 56
    - trap, 58
  - integer load, store, and cache management instructions, exception priorities for, 196
  - integer processing, 69
  - interfaces
    - interrupt controller, 36
  - internal debug mode, 213
  - interrupt
    - alignment interrupt, 179
    - data storage interrupt, 175
    - external input interrupt, 179
    - instruction
      - partially executed, 158
    - Instruction Storage, 178
    - instruction storage interrupt, 178
    - instruction TLB miss interrupt, 188
    - machine check interrupt, 172
    - masking, 192
      - guidelines for system software, 194
    - ordering, 192, 194
      - guidelines for system software, 194
    - program interrupt, 180
      - illegal instruction exception, 181
      - privileged instruction exception, 181
      - trap exception, 183
    - system call interrupt, 184
    - type
      - Alignment, 179
      - Auxiliary Processor Unavailable, 184
      - Critical Input, 172
      - Data Storage, 175

- Data TLB Error, 186
  - Debug, 188
  - Decrementer, 185
  - External Input, 179
  - Fixed-Interval Timer, 185
  - Floating-Point Unavailable, 183
  - Instruction TLB Error, 188
  - Machine Check, 172
  - Program interrupt, 180
  - System Call, 184
  - Watchdog Timer, 186
  - interrupt (IRPT) debug events, 230
  - interrupt and exception handling registers
    - ESR, 166
  - interrupt classes
    - asynchronous, 153
    - critical and non-critical, 155
    - machine check, 155
    - synchronous, 153
  - interrupt controller interface, 36
  - interrupt processing, 156
    - interrupt vector, 156
  - interrupt vector, 156
  - Interrupts, 153
  - interrupts
    - definitions, 169
    - imprecise, 154
    - order, 194
    - ordering and masking, 192
    - ordering and software, 193
    - partially executed instructions, 158
    - precise, 154
    - registers, processing, 159
    - synchronous and imprecise, 154
    - synchronous and precise, 154
    - types
      - alignment, 179
      - auxiliary processor unavailable, 184
      - data storage, 175
      - data TLB error, 186
      - debug, 188
      - decrementer, 185
      - definitions, 169
      - external inputs, 179
      - fixed interval timer, 185
      - floating point unavailable, 183
      - instruction storage, 178
      - instruction TLB error, 188
      - machine check, 172
      - program, 180
      - watchdog timer, 186
    - vectors, 156
  - INV0–INV3, 482
  - isel, 310
  - isync, 311
  - ITV0–ITV3, 483
  - IVLIM, 484
  - IVOR0–IVOR15, 485
  - IVPR, 486
- ## L
- lbz, 312
  - lbzu, 313
  - lbzx, 315
  - lha, 316
  - lhau, 317
  - lhax, 319
  - lhbrx, 320
  - lhz, 321
  - lhzu, 322
  - lhzux, 323
  - lhzx, 324
  - li, 251
  - lis, 254
  - little endian
    - structure mapping, 44
  - little endian mapping, 44
  - little endian, defined, 43
  - lmw, 325
  - load and store alignment, 114
  - load operations, 115
  - locking, cache lines, 96
  - logical compare, 69
  - LR, 65, 487
  - lswi, 326
  - lswx, 328
  - lwarx, 330
  - lwz, 332
  - lwzu, 333
  - lwzux, 334
  - lwzx, 335
- ## M
- M storage attribute, 141
  - macchw, 336
  - macchws, 337
  - macchwsu, 338
  - macchwu, 339
  - machhw, 340
  - machhwsu, 342
  - machhwu, 343
  - Machine Check, 155
  - Machine Check interrupt, 172
  - machine check interrupts, 155, 172
  - Machine State Register. *See also* MSR
  - maclhw, 344
  - maclhws, 345, 382
  - maclhwu, 347

masking and ordering interrupts, 192  
 mbar, 348  
 mcrf, 350  
 mcrxr, 351  
 MCSR, 488  
 MCSRR0, 489  
 MCSRR1, 490  
 memory coherence required, 141  
 memory management unit, 32  
 memory management. *See also* MMU  
 memory map, 39  
 memory organization, 39  
 mfcrr, 352  
 mfdcr, 353  
 mfmsr, 159, 354  
 mfspr, 355  
 MMU  
   change status management, 150  
   overview, 129  
   page reference, 150  
   PowerPC Book-E MMU Architecture, nonsupported  
     features, 130  
   support for Power PC Book-E MMU architecture, 129  
   TLB management instructions  
     overview, 148  
     read/write (tlbre, tlbwe), 149  
     search (tlbsx), 148  
 MMUCR, 143, 491  
 mr, 384  
 mr., 384  
 MSR, 159, 492  
   defined, 52  
 msync, 358  
 mtcr, 359  
 mtcrf, 359  
 mtdcr, 360  
 mtmsr, 159  
 mtspr, 362  
 mulchw, 365  
 mulchwu, 366  
 mulhhw, 367  
 mulhhwu, 368  
 mulhwu, 370  
 mulhwu., 370  
 mullhw, 371  
 mullhwu, 372  
 mulli, 373  
 mullw, 374  
 mullw., 374  
 mullwo, 374  
 mullwo., 374

## N

nand, 375

nand., 375  
 neg, 376  
 neg., 376  
 nego, 376  
 nego., 376  
 nmacchw, 377  
 nmacchws, 378  
 nmachhw, 379  
 nmachhws, 380  
 nmaclhw, 381  
 nmaclhws, 382  
 non-critical interrupts, 155  
 nop, 386  
 nor, 383  
 nor., 383  
 not, 383  
 not., 383  
 notation, 25, 245, 508  
 notational conventions, 25

## O

opcodes, 544  
   allocated instruction, 543  
   preserved instruction, 543  
 operands  
   storage, 39  
 operations  
   DCC, 113  
   ICC, 101  
   line flush, 117  
   load, 115  
   store, 115  
 or, 384  
 or., 384  
 orc, 385  
 orc., 385  
 ordering  
   storage access, 120  
 ordering and masking interrupts, 192  
 ori, 386  
 oris, 387

## P

page management, 150  
 partially executed instructions, 158  
 PID, 146, 494  
 PIR, 74, 495  
 portability, instruction set, 244  
 precise interrupts, 154  
 prefetch mechanism, speculative, 102  
 preserved instructions, exception priorities for, 200  
 primary opcodes, 544

priorities, exception, 195  
 privileged instructions, 78  
 privileged mode, 78  
 privileged operation, 78  
 privileged SPRs, 79  
 problem state, 78  
 processor control instruction summary, 59  
 processor control instructions  
   CR logical, 60  
   register management, 60  
   synchronization, 60  
   system linkage, 60  
 processor control registers, 72  
 Program interrupt, 180  
 program interrupts, 180  
 pseudocode, 245  
 PVR, 73, 496

## R

R0-R31, 477  
 reading the time base, 204  
 register

  CSRR0, 162, 163  
   CSRR1, 162, 163  
   ESR, 166  
   PID, 146  
   SRR0, 161  
   SRR1, 161

registers, 47

  branching control, 65  
   CCR0, 74, 105, 107, 122, 448  
   CCR1, 450  
   CR, 52, 66, 452  
   CSRR0, 453  
   CSRR1, 454  
   CTR, 66, 455  
   DAC1–DAC2, 240, 456  
   DBCR0, 233, 237, 457  
   DBCR1, 234, 459  
   DBCR2, 461  
   DBDR, 241  
   DBSR, 238  
   DCDBTRH, 123, 466  
   DCDBTRL, 123, 467  
   DCR numbering, 445  
   DEAR, 468  
   DEC, 205, 469  
   DECAR, 206, 470  
   DNV0–DNV3, 471  
   DTV0–DTV3, 472  
   DVC1–DVC2, 240  
   DVLIM, 474  
   ESR, 166, 475  
   GPR0–GPR31, 477

  GPRs, 69  
   IAC1–IAC4, 473, 478  
   IAC1–IAC4, 239  
   ICDBDR, 479  
   ICDBTRH, 480  
   ICDBTRL, 481  
   interrupt processing, 159  
   INV0–INV3, 482  
   ITV0–ITV3, 483  
   IVLIM, 484  
   IVOR0–IVOR15, 485  
   IVPR, 486  
   LR, 65, 487  
   MCSR, 488  
   MCSRR0, 489  
   MCSRR1, 490  
   MMUCR, 143, 491  
   MSR, 52, 159, 492  
   PID, 494  
   PIR, 74, 495  
   processor control, 72  
   PVR, 73, 496  
   R0-R31, 477  
   RSTCFG, 77, 497  
   SPRG0–SPRG7, 73  
   SPRG0–SPRG3, 73  
   SPRG0–SPRG7, 498  
   SRR0, 499  
   SRR1, 500  
   storage control, 143  
   TBL, 501  
   TBU, 502  
   TCR, 206, 207, 209, 503  
   TSR, 208, 210, 504  
   types, 51  
     CR, 52  
     DCR, 52  
     GPR, 51  
     MSR, 52  
     SPR, 51  
   USPRG0, 73, 505  
   XER, 70, 506  
 registers, device control, 52  
 registers, summary, 47  
 replacement policy, cache line, 94  
 requirements  
   software  
     interrupt ordering, 193  
 reservation bit, 330, 414  
 reserved instructions, exception priorities for, 201  
 reserved-illegal instructions, 544  
 reserved-nop instructions, 544  
 reset  
   debug, 232  
 return (RET) debug events, 229  
 return from interrupt instructions, exception priorities for,



200  
 rfc1, 388  
 rfi, 161, 389  
 rfmci, 390  
 rlwimi, 391  
 rlwimi., 391  
 rlwinm, 392  
 rlwinm., 392  
 rlwnm, 394  
 rlwnm., 394  
 rotlw, 394  
 rotlw., 394  
 rotlwi, 393  
 rotlwi., 393  
 rotrwi, 393  
 rotrwi., 393  
 RSTCFG, 77, 497

## S

Save/Restore Register 0, 161  
 Save/Restore Register 1, 161  
 sc, 395  
 secondary opcodes, 544  
 self-modifying code, 103  
 shadow TLB arrays, 147  
 slw, 396  
 slw., 396  
 slwi, 393  
 slwi., 393  
 software  
     interrupt ordering requirements, 193  
 Special Purpose Registers. *See also* SPRs  
 speculative fetching, 79  
 speculative prefetch mechanism, 102  
 SPR0-SPRG7, 73  
 SPRG0-SPRG3, 73  
 SPRG0-SPRG7, 498  
 SPRs  
     defined, 51  
 srw, 397  
 srw., 397  
 srwi, 398  
 srwi., 398  
 SRR0, 161, 499  
 SRR1, 161, 500  
 srw, 399  
 srw., 399  
 srwi, 393  
 srwi., 393  
 stb, 400  
 stbu, 401  
 stbux, 402  
 stbx, 403  
 sth, 404

sthbrx, 405  
 sthu, 406  
 sthux, 407  
 sthx, 408  
 stmw, 409  
 storage access ordering, 120  
 storage attributes  
     caching inhibited, 141  
     endian, 142  
     guarded, 141  
     Memory Coherence Required, 141  
     supported combinations, 143  
     user-definable (U0–U3), 142  
     write-through required, 141  
 storage control instruction summary, 61  
 storage control instructions  
     cache management, 61  
     TLB management, 61  
 storage operands, 39  
 storage synchronization, 81  
 storage synchronization instruction summary, 61  
 store gathering, 116  
 store miss  
     allocation of data cache line, 116  
 store operations, 115  
 structure mapping  
     big endian, 44  
     little endian, 44  
 stswi, 409  
 stw, 412  
 stwbrx, 413  
 stwcx., 414  
 stwu, 416  
 stwux, 417  
 stwx, 418  
 sub, 419  
 sub., 419  
 subc, 420  
 subc., 420  
 subco, 420  
 subco., 420  
 subf, 419  
 subf., 419  
 subfc, 420  
 subfc., 420  
 subfco, 420  
 subfco., 420  
 subfe, 421  
 subfe., 421  
 subfeo, 421  
 subfeo., 421  
 subfic, 422  
 subfme, 423  
 subfme., 423  
 subfmeo, 423  
 subfmeo., 423

- subfo, 419
- subfo., 419
- subfze, 424
- subfze., 424
- subfzeo, 424
- subfzeo., 424
- subi, 251
- subic, 252
- subic., 253
- subis, 254
- subo, 419
- subo., 419
- supervisor state, 78
- synchronization
  - architectural references, 79
  - context, 80
  - execution, 81
  - storage, 81
- synchronous interrupt class, 153
- synonyms, instruction cache, 104
- system call instruction, exception priorities for, 200
- System Call interrupt, 184

## T

- TBL, 501
- TBU, 502
- TCR, 207, 209, 503
- time base
  - defined, 204
  - reading, 204
  - writing, 205
- timer freeze (debug), 232
- timers
  - DEC, 205
  - DECAR, 206
  - decrementer, 205, 206
  - FIT, 206
  - fixed interval timer, 206
  - freezing the timer facilities, 211
  - TCR, 209
  - TSR, 210
  - watchdog timer, 207
  - watchdog timer state machine, 209
- TLB
  - entry fields
    - E, 133
    - EPN, 131
    - ERP, 132
    - G, 133
    - I, 132
    - M, 133
    - RPN, 132
    - SIZE, 131
    - TID, 131

- TS, 131
- U0, 132
- U1, 132
- U2, 132
- U3, 132
- UR, 133
- UW, 133
- UX, 133
- V, 131
- W, 132
- overview, 130
- shadow arrays, 147
- TLB management instructions
  - overview, 148
- tlbre, 425
- tlbsx, 427
- tlbsx., 427
- tlbsync, 428
- tlbwe, 429
- trace debug mode, 215
- transient mechanism, cache, 96
- translation lookaside buffer. *See also* TLB
- trap, 431
- trap (TRAP) debug events, 228
- trap instructions
  - exception priorities for, 199
- TSR, 208, 210, 504
- tw, 430
- tweq, 431
- tweqi, 434
- twge, 431
- twgei, 434
- twgle, 431
- twgt, 431
- twgti, 434
- twi, 433
- twle, 431
- twlei, 434
- twlgei, 434
- twlgt, 431
- twlgti, 434
- twlle, 431
- twllei, 434
- twllt, 431
- twllti, 434
- twlng, 431
- twlngi, 434
- twlnl, 431
- twlnli, 434
- twlt, 431
- twlti, 434
- twne, 431
- twnei, 434
- twng, 431
- twngi, 434
- twnl, 432

twlli, 434

## U, V, W

U0–U3 storage attributes, 142  
unconditional (UDE) debug events, 231  
units  
    memory management, 32  
user mode, 78  
USPRG0, 73, 505  
W storage attribute, 141  
Watchdog Timer interrupt, 186  
watchdog timer interrupts, 186  
write-through required, 141  
writing the time base, 205  
wrtee, 435  
wrteei, 436

## X

XER, 70, 506  
    carry (CA) field, 72  
    overflow (OV) field, 71  
    summary overflow (SO) field, 71  
    transfer byte count (TBC) field, 72  
xor, 437  
xori, 438



## Revision Log

Revision Date	Contents of Modification
07/15/03	Updated release of the manual (SA14-2613-03) after test build review
06/30/03	Test build of the manual for review
06/24/03	Minor updates since the previous release. Template related fixes to all files Disclaimer page updated
10/18/2002	Reformatted figures and tables to division standard template, no content revisions.
9/12/2002	Content and format revisions summarized by chapter: Ch. 2: CCR1 updated Ch. 4: CCR0 and CCR1 updated. Sections on data cache parity insertion and simulating parity errors revised. Ch. 6: Revised MCSR, MCSRR0, MCSRR1, and descriptions of Machine Check Interrupt handling. Ch. 9: Added extended mnemonics to mfspr, mtspr for MCSR, MCSRR0, MCSRR1 Ch10: Fixed format problems in MCSR, MCSRR0, MCSRR1, MSR, SRR1 Appendix A: Added extended mnemonics to mfspr, mtspr for MCSR, MCSRR0, MCSRR1
7/25/2002	Reformatted to division standard template, no content revisions.

