



PowerPC 405
Embedded Processor Core
User's Manual

Fifth Edition (December 2001)

This edition of *IBM PPC405 Embedded Processor Core User's Manual* applies to the IBM PPC405 32-bit embedded processor core, until otherwise indicated in new versions or application notes.

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS MANUAL "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

IBM does not warrant that the products in this publication, whether individually or as one or more groups, will meet your requirements or that the publication or the accompanying product descriptions are error-free.

This publication could contain technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time.

It is possible that this publication may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country. Any reference to an IBM licensed program in this publication is not intended to state or imply that you can use only IBM's licensed program. You can use any functionally equivalent program instead.

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the written permission of IBM.

Requests for copies of this publication and for technical information about IBM products should be made to your IBM Authorized Dealer or your IBM Marketing Representative.

Address technical queries about this product to ppcsupp@us.ibm.com

Address comments about this publication to:

IBM Corporation
Department YM5A
P.O. Box 12195
Research Triangle Park, NC 27709

IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1996, 2001. All rights reserved

4 3 2 1

Notice to U.S. Government Users – Documentation Related to Restricted Rights – Use, duplication, or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corporation.

Patents and Trademarks

IBM may have patents or pending patent applications covering the subject matter in this publication. The furnishing of this publication does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 208 Harbor Drive, Stamford, CT 06904, United States of America.

The following terms are trademarks of IBM Corporation:

IBM

PowerPC

PowerPC Architecture

PowerPC Embedded Controllers

RISCWatch

Other terms which are trademarks are the property of their respective owners.

Contents

Figures	xv
Tables	xviii
About This Book	xxi
Who Should Use This Book	xxi
How to Use This Book	xxi
Conventions	xxii
Chapter 1. Overview	1-1
PPC405 Features	1-1
PowerPC Architecture	1-3
The PPC405 as a PowerPC Implementation	1-3
Processor Core Organization	1-4
Instruction and Data Cache Controllers	1-4
Instruction Cache Unit	1-4
Data Cache Unit	1-5
Memory Management Unit	1-5
Timer Facilities	1-6
Debug	1-7
Development Tool Support	1-7
Debug Modes	1-7
Core Interfaces	1-7
Processor Local Bus	1-8
Device Control Register Bus	1-8
Clock and Power Management	1-8
JTAG	1-8
Interrupts	1-8
Auxiliary Processor Unit	1-8
On-Chip Memory	1-8
Data Types	1-8
Processor Core Register Set Summary	1-9
General Purpose Registers	1-9
Special Purpose Registers	1-9
Machine State Register	1-9
Condition Register	1-9
Device Control Registers	1-9
Addressing Modes	1-10
Chapter 2. Programming Model	2-1
User and Privileged Programming Models	2-1
Memory Organization and Addressing	2-1
Storage Attributes	2-2
Registers	2-2
General Purpose Registers (R0-R31)	2-5
Special Purpose Registers	2-5
Count Register (CTR)	2-6
Link Register (LR)	2-7
Fixed Point Exception Register (XER)	2-7
Special Purpose Register General (SPRG0–SPRG7)	2-9
Processor Version Register (PVR)	2-10
Condition Register (CR)	2-10
CR Fields after Compare Instructions	2-11

The CR0 Field	2-12
The Time Base	2-13
Machine State Register (MSR)	2-13
Device Control Registers	2-15
Data Types and Alignment	2-16
Alignment for Storage Reference and Cache Control Instructions	2-16
Alignment and Endian Operation	2-17
Summary of Instructions Causing Alignment Exceptions	2-17
Byte Ordering	2-17
Structure Mapping Examples	2-18
Big Endian Mapping	2-19
Little Endian Mapping	2-19
Support for Little Endian Byte Ordering	2-19
Endian (E) Storage Attribute	2-19
Fetching Instructions from Little Endian Storage Regions	2-20
Accessing Data in Little Endian Storage Regions	2-21
PowerPC Byte-Reverse Instructions	2-21
Instruction Processing	2-23
Branch Processing	2-24
Unconditional Branch Target Addressing Options	2-24
Conditional Branch Target Addressing Options	2-24
Conditional Branch Condition Register Testing	2-25
BO Field on Conditional Branches	2-25
Branch Prediction	2-26
Speculative Accesses	2-27
Speculative Accesses in the PPC405	2-27
Prefetch Distance Down an Unresolved Branch Path	2-28
Prefetch of Branches to the CTR and Branches to the LR	2-28
Preventing Inappropriate Speculative Accesses	2-28
Fetching Past an Interrupt-Causing or Interrupt-Returning Instruction	2-28
Fetching Past tw or twi Instructions	2-29
Fetching Past an Unconditional Branch	2-29
Suggested Locations of Memory-Mapped Hardware	2-29
Summary	2-30
Privileged Mode Operation	2-30
MSR Bits and Exception Handling	2-31
Privileged Instructions	2-31
Privileged SPRs	2-32
Privileged DCRs	2-32
Synchronization	2-33
Context Synchronization	2-33
Execution Synchronization	2-35
Storage Synchronization	2-35
Instruction Set	2-36
Instructions Specific to the IBM PowerPC Embedded Environment	2-37
Storage Reference Instructions	2-37
Arithmetic Instructions	2-38
Logical Instructions	2-39
Compare Instructions	2-39
Branch Instructions	2-40
CR Logical Instructions	2-40
Rotate Instructions	2-40
Shift Instructions	2-41
Cache Management Instructions	2-41
Interrupt Control Instructions	2-41
TLB Management Instructions	2-42

Processor Management Instructions	2-42
Extended Mnemonics	2-42
Chapter 3. Initialization	3-1
Processor State After Reset	3-1
Machine State Register Contents after Reset	3-2
Contents of Special Purpose Registers after Reset	3-3
PPC405 Initial Processor Sequencing	3-3
Initialization Requirements	3-4
Initialization Code Example	3-5
Chapter 4. Cache Operations	4-1
ICU and DCU Organization and Sizes	4-2
ICU Overview	4-3
ICU Operations	4-4
Instruction Cachability Control	4-5
Instruction Cache Synonyms	4-5
ICU Coherency	4-6
DCU Overview	4-6
DCU Operations	4-6
DCU Write Strategies	4-7
DCU Load and Store Strategies	4-8
Data Cachability Control	4-8
DCU Coherency	4-9
Cache Instructions	4-9
ICU Instructions	4-9
DCU Instructions	4-10
Cache Control and Debugging Features	4-11
CCR0 Programming Guidelines	4-13
ICU Debugging	4-14
DCU Debugging	4-15
DCU Performance	4-16
Pipeline Stalls	4-16
Cache Operation Priorities	4-17
Simultaneous Cache Operations	4-17
Sequential Cache Operations	4-18
Chapter 5. Fixed-Point Interrupts and Exceptions	5-1
Architectural Definitions and Behavior	5-1
Behavior of the PPC405 Processor Core Implementation	5-2
Interrupt Handling Priorities	5-3
Critical and Noncritical Interrupts	5-5
General Interrupt Handling Registers	5-7
Machine State Register (MSR)	5-7
Save/Restore Registers 0 and 1 (SRR0–SRR1)	5-9
Save/Restore Registers 2 and 3 (SRR2–SRR3)	5-9
Exception Vector Prefix Register (EVPR)	5-10
Exception Syndrome Register (ESR)	5-11
Data Exception Address Register (DEAR)	5-13
Critical Input Interrupts	5-13
Machine Check Interrupts	5-14
Instruction Machine Check Handling	5-14
Data Machine Check Handling	5-15
Data Storage Interrupt	5-16
Instruction Storage Interrupt	5-17
External Interrupt	5-18
External Interrupt Handling	5-18

Alignment Interrupt	5-19
Program Interrupt	5-20
FPU Unavailable Interrupt	5-21
System Call Interrupt	5-22
APU Unavailable Interrupt	5-22
Programmable Interval Timer (PIT) Interrupt	5-22
Fixed Interval Timer (FIT) Interrupt	5-23
Watchdog Timer Interrupt	5-24
Data TLB Miss Interrupt	5-25
Instruction TLB Miss Interrupt	5-25
Debug Interrupt	5-26
Chapter 6. Timer Facilities	6-1
Time Base	6-1
Reading the Time Base	6-3
Writing the Time Base	6-3
Programmable Interval Timer (PIT)	6-4
Fixed Interval Timer (FIT)	6-5
Watchdog Timer	6-6
Timer Status Register (TSR)	6-8
Timer Control Register (TCR)	6-9
Chapter 7. Memory Management	7-1
MMU Overview	7-1
Address Translation	7-1
Translation Lookaside Buffer (TLB)	7-2
Unified TLB	7-2
TLB Fields	7-3
Page Identification Fields	7-3
Translation Field	7-4
Access Control Fields	7-5
Storage Attribute Fields	7-5
Shadow Instruction TLB	7-6
ITLB Accesses	7-7
Shadow Data TLB	7-7
DTLB Accesses	7-7
Shadow TLB Consistency	7-7
TLB-Related Interrupts	7-9
Data Storage Interrupt	7-10
Instruction Storage Interrupt	7-10
Data TLB Miss Interrupt	7-11
Instruction TLB Miss Interrupt	7-11
Program Interrupt	7-11
TLB Management	7-11
TLB Search Instructions (tlbsx/tlbsx.)	7-12
TLB Read/Write Instructions (tlbre/tlbwe)	7-12
TLB Invalidate Instruction (tlbia)	7-12
TLB Sync Instruction (tlbsync)	7-12
Recording Page References and Changes	7-12
Access Protection	7-13
Access Protection Mechanisms in the TLB	7-13
General Access Protection	7-13
Execute Permissions	7-14
Write Permissions	7-14
Zone Protection	7-14
Access Protection for Cache Control Instructions	7-16
Access Protection for String Instructions	7-17

Real-Mode Storage Attribute Control	7-17
Storage Attribute Control Registers	7-19
Data Cache Write-through Register (DCWR)	7-19
Data Cache Cachability Register (DCCR)	7-20
Instruction Cache Cachability Register (ICCR)	7-20
Storage Guarded Register (SGR)	7-20
Storage User-defined 0 Register (SU0R)	7-20
Storage Little-Endian Register (SLER)	7-20
Chapter 8. Debugging	8-1
Development Tool Support	8-1
Debug Modes	8-1
Internal Debug Mode	8-1
External Debug Mode	8-2
Debug Wait Mode	8-2
Real-time Trace Debug Mode	8-3
Processor Control	8-3
Processor Status	8-4
Debug Registers	8-4
Debug Control Registers	8-4
Debug Control Register 0 (DBCR0)	8-4
Debug Control Register1 (DBCR1)	8-6
Debug Status Register (DBSR)	8-7
Instruction Address Compare Registers (IAC1–IAC4)	8-9
Data Address Compare Registers (DAC1–DAC2)	8-9
Data Value Compare Registers (DVC1–DVC2)	8-10
Debug Events	8-10
Instruction Complete Debug Event	8-11
Branch Taken Debug Event	8-11
Exception Taken Debug Event	8-11
Trap Taken Debug Event	8-12
Unconditional Debug Event	8-12
IAC Debug Event	8-12
IAC Exact Address Compare	8-12
IAC Range Address Compare	8-12
DAC Debug Event	8-13
DAC Exact Address Compare	8-13
DAC Range Address Compare	8-14
DAC Applied to Cache Instructions	8-15
DAC Applied to String Instructions	8-16
Data Value Compare Debug Event	8-16
Imprecise Debug Event	8-19
Debug Interface	8-19
IEEE 1149.1 Test Access Port (JTAG Debug Port)	8-19
JTAG Connector	8-20
JTAG Instructions	8-21
JTAG Boundary Scan	8-21
Trace Port	8-22
Chapter 9. Instruction Set	9-1
Instruction Set Portability	9-1
Instruction Formats	9-2
Pseudocode	9-2
Operator Precedence	9-5
Register Usage	9-5
Alphabetical Instruction Listing	9-5
add	9-6

addc	9-7
adde	9-8
addi	9-9
addic	9-10
addic.	9-11
addis	9-12
addme	9-13
addze	9-14
and	9-15
andc	9-16
andi.	9-17
andis.	9-18
b	9-19
bc	9-20
bcctr	9-26
bclr	9-30
cmp	9-34
cmpi	9-35
cmpl	9-36
cmpli	9-37
cntlzw	9-38
crand	9-39
crandc	9-40
creqv	9-41
crnand	9-42
crnor	9-43
cror	9-44
crorc	9-45
crxor	9-46
dcba	9-47
dcbf	9-49
dcbi	9-50
dcbst	9-51
dcbt	9-52
dcbstst	9-53
dcbz	9-54
dccci	9-56
dcread	9-57
divw	9-59
divwu	9-60
eieio	9-61
eqv	9-62
extsb	9-63
extsh	9-64
icbi	9-65
icbt	9-66
iccci	9-67
icread	9-68
isync	9-70
lbz	9-71
lbzu	9-72
lbzux	9-73
lbzx	9-74
lha	9-75

lhau	9-76
lhaux	9-77
lhax	9-78
lhbrx	9-79
lhz	9-80
lhzu	9-81
lhzux	9-82
lhzx	9-83
lmw	9-84
lswi	9-85
lswx	9-87
lwarx	9-89
lwbrx	9-90
lwz	9-91
lwzu	9-92
lwzux	9-93
lwzx	9-94
macchw	9-95
macchws	9-96
macchwsu	9-97
macchwu	9-98
machhw	9-99
machhws	9-100
machhwsu	9-101
machhwu	9-102
maclhw	9-103
maclhws	9-104
maclhwsu	9-105
maclhwu	9-106
mcrf	9-107
mcrxr	9-108
mfcf	9-109
mfdcf	9-110
mfmsr	9-111
mfspr	9-112
mftb	9-114
mtcrf	9-116
mtdcf	9-117
mtmsr	9-118
mtspr	9-119
mulchw	9-121
mulchwu	9-122
mulhhw	9-123
mulhhwu	9-124
mulhw	9-125
mulhwu	9-126
mullhw	9-127
mullhwu	9-128
mulli	9-129
mullw	9-130
nand	9-131
neg	9-132
nmacchw	9-133
nmacchws	9-134

nmachw	9-135
nmachws	9-136
nmaclhw	9-137
nmaclhws	9-138
nor	9-139
or	9-140
orc	9-141
ori	9-142
oris	9-143
rfci	9-144
rfi	9-145
rlwimi	9-146
rlwinm	9-147
rlwnm	9-150
sc	9-151
slw	9-152
sraw	9-153
srawi	9-154
srw	9-155
stb	9-156
stbu	9-157
stbux	9-158
stbx	9-159
sth	9-160
sthbrx	9-161
sthu	9-162
sthux	9-163
sthx	9-164
stmw	9-165
stswi	9-166
stswx	9-167
stw	9-169
stwbrx	9-170
stwcx	9-171
stwu	9-173
stwux	9-174
stwx	9-175
subf	9-176
subfc	9-177
subfe	9-178
subfic	9-179
subfme	9-180
subfze	9-181
sync	9-182
tlbia	9-183
tlbre	9-184
tlbsx	9-186
tlbsync	9-187
tlbwe	9-188
tw	9-190
twi	9-193
wrtee	9-196
wrteei	9-197
xor	9-198

xori	9-199
xoris	9-200
Chapter 10. Register Summary	10-1
Reserved Registers	10-1
Reserved Fields	10-1
General Purpose Registers	10-1
Machine State Register and Condition Register	10-1
Special Purpose Registers	10-2
Time Base Registers	10-4
Device Control Registers	10-4
Alphabetical Listing of PPC405 Registers	10-5
CCR0	10-6
CR	10-8
CTR	10-9
DAC1–DAC2	10-10
DBCR0	10-11
DBCR1	10-13
DBSR	10-15
DCCR	10-17
DCWR	10-19
DEAR	10-21
DVCR1–DVCR2	10-22
ESR	10-23
EVPR	10-25
GPR0–GPR31	10-26
IAC1–IAC4	10-27
ICCR	10-28
ICDBDR	10-30
LR	10-31
MSR	10-32
PID	10-34
PIT	10-35
PVR	10-36
SGR	10-37
SLER	10-39
SPRG0–SPRG7	10-41
SRR0	10-42
SRR1	10-43
SRR2	10-44
SRR3	10-45
SU0R	10-46
TBL	10-48
TBU	10-49
TCR	10-50
TSR	10-51
USPRG0	10-52
XER	10-53
ZPR	10-54
A. Instruction Summary	A-1
Instruction Set and Extended Mnemonics – Alphabetical	A-1
Instructions Sorted by Opcode	A-33
Instruction Formats	A-41
Instruction Fields	A-41

Instruction Format Diagrams	A-43
I-Form A-44	
B-Form A-44	
SC-Form A-44	
D-Form A-44	
X-Form A-45	
XL-Form A-45	
XFX-Form A-46	
X0-Form A-46	
M-Form A-46	
B. Instructions by Category	B-1
Implementation-Specific Instructions	B-1
Instructions in the IBM PowerPC Embedded Environment	B-5
Privileged Instructions	B-7
Assembler Extended Mnemonics	B-9
Storage Reference Instructions	B-29
Arithmetic and Logical Instructions	B-33
Condition Register Logical Instructions	B-37
Branch Instructions	B-38
Comparison Instructions	B-39
Rotate and Shift Instructions	B-40
Cache Control Instructions	B-41
Interrupt Control Instructions	B-42
TLB Management Instructions	B-42
Processor Management Instructions	B-44
C. Code Optimization and Instruction Timings	C-1
Code Optimization Guidelines	C-1
Condition Register Bits for Boolean Variables	C-1
CR Logical Instruction for Compound Branches	C-1
Floating-Point Emulation	C-1
Cache Usage	C-2
CR Dependencies	C-2
Branch Prediction	C-2
Alignment	C-2
Instruction Timings	C-3
General Rules	C-3
Branches	C-3
Multiplies	C-4
Scalar Load Instructions	C-5
Scalar Store Instructions	C-6
Alignment in Scalar Load and Store Instructions	C-6
String and Multiple Instructions	C-6
Loads and Store Misses	C-7
Instruction Cache Misses	C-7
Index	X-1

Figures

Figure 1-1. PPC405 Block Diagram	1-4
Figure 2-1. PPC405 Programming Model—Registers	2-4
Figure 2-2. General Purpose Registers (R0-R31)	2-5
Figure 2-3. Count Register (CTR)	2-7
Figure 2-4. Link Register (LR)	2-7
Figure 2-5. Fixed Point Exception Register (XER)	2-8
Figure 2-6. Special Purpose Register General (SPRG0–SPRG7)	2-10
Figure 2-7. Processor Version Register (PVR)	2-10
Figure 2-8. Condition Register (CR)	2-11
Figure 2-9. Machine State Register (MSR)	2-14
Figure 2-10. PPC405 Data Types	2-16
Figure 2-11. Normal Word Load or Store (Big Endian Storage Region)	2-22
Figure 2-12. Byte-Reverse Word Load or Store (Little Endian Storage Region)	2-22
Figure 2-13. Byte-Reverse Word Load or Store (Big Endian Storage Region)	2-22
Figure 2-14. Normal Word Load or Store (Little Endian Storage Region)	2-23
Figure 2-15. PPC405 Instruction Pipeline	2-24
Figure 4-1. Instruction Flow	4-4
Figure 4-2. Core Configuration Register 0 (CCR0)	4-11
Figure 4-3. Instruction Cache Debug Data Register (ICDBDR)	4-14
Figure 5-1. Machine State Register (MSR)	5-7
Figure 5-2. Save/Restore Register 0 (SRR0)	5-9
Figure 5-3. Save/Restore Register 1 (SRR1)	5-9
Figure 5-4. Save/Restore Register 2 (SRR2)	5-10
Figure 5-5. Save/Restore Register 3 (SRR3)	5-10
Figure 5-6. Exception Vector Prefix Register (EVPR)	5-11
Figure 5-7. Exception Syndrome Register (ESR)	5-11
Figure 5-8. Data Exception Address Register (DEAR)	5-13
Figure 6-1. Relationship of Timer Facilities to the Time Base	6-1
Figure 6-2. Time Base Lower (TBL)	6-2
Figure 6-3. Time Base Upper (TBU)	6-2
Figure 6-4. Programmable Interval Timer (PIT)	6-5
Figure 6-5. Watchdog Timer State Machine	6-7
Figure 6-6. Timer Status Register (TSR)	6-8
Figure 6-7. Timer Control Register (TCR)	6-9
Figure 7-1. Effective to Real Address Translation Flow	7-2
Figure 7-2. TLB Entries	7-3
Figure 7-3. ITLB/DTLB/UTLB Address Resolution	7-9
Figure 7-4. Process ID (PID)	7-14
Figure 7-5. Zone Protection Register (ZPR)	7-15
Figure 7-6. Generic Storage Attribute Control Register	7-19
Figure 8-1. Debug Control Register 0 (DBCR0)	8-4
Figure 8-2. Debug Control Register 1 (DPCR1)	8-6

Figure 8-3. Debug Status Register (DBSR)	8-8
Figure 8-4. Instruction Address Compare Registers (IAC1–IAC4)	8-9
Figure 8-5. Data Address Compare Registers (DAC1–DAC2)	8-10
Figure 8-6. Data Value Compare Registers (DVC1–DVC2)	8-10
Figure 8-7. Inclusive IAC Range Address Compares	8-13
Figure 8-8. Exclusive IAC Range Address Compares	8-13
Figure 8-9. Inclusive DAC Range Address Compares	8-15
Figure 8-10. Exclusive DAC Range Address Compares	8-15
Figure 8-11. JTAG Connector Physical Layout (Top View)	8-20
Figure 10-1. Core Configuration Register 0 (CCR0)	10-6
Figure 10-2. Condition Register (CR)	10-8
Figure 10-3. Count Register (CTR)	10-9
Figure 10-4. Data Address Compare Registers (DAC1–DAC2)	10-10
Figure 10-5. Debug Control Register 0 (DBCR0)	10-11
Figure 10-6. Debug Control Register 1 (DBCR1)	10-13
Figure 10-7. Debug Status Register (DBSR)	10-15
Figure 10-8. Data Cache Cachability Register (DCCR)	10-17
Figure 10-9. Data Cache Write-through Register (DCWR)	10-19
Figure 10-10. Data Exception Address Register (DEAR)	10-21
Figure 10-11. Data Value Compare Registers (DVC1–DVC2)	10-22
Figure 10-12. Exception Syndrome Register (ESR)	10-23
Figure 10-13. Exception Vector Prefix Register (EVPR)	10-25
Figure 10-14. General Purpose Registers (R0-R31)	10-26
Figure 10-15. Instruction Address Compare Registers (IAC1–IAC4)	10-27
Figure 10-16. Instruction Cache Cachability Register (ICCR)	10-28
Figure 10-17. Instruction Cache Debug Data Register (ICDBDR)	10-30
Figure 10-18. Link Register (LR)	10-31
Figure 10-19. Machine State Register (MSR)	10-32
Figure 10-20. Process ID (PID)	10-34
Figure 10-21. Programmable Interval Timer (PIT)	10-35
Figure 10-22. Processor Version Register (PVR)	10-36
Figure 10-23. Storage Guarded Register (SGR)	10-37
Figure 10-24. Storage Little-Endian Register (SLER)	10-39
Figure 10-25. Special Purpose Registers General (SPRG0–SPRG7)	10-41
Figure 10-26. Save/Restore Register 0 (SRR0)	10-42
Figure 10-27. Save/Restore Register 1 (SRR1)	10-43
Figure 10-28. Save/Restore Register 2 (SRR2)	10-44
Figure 10-29. Save/Restore Register 3 (SRR3)	10-45
Figure 10-30. Storage User-defined 0 Register (SU0R)	10-46
Figure 10-31. Time Base Lower (TBL)	10-48
Figure 10-32. Time Base Upper (TBU)	10-49
Figure 10-33. Timer Control Register (TCR)	10-50
Figure 10-34. Timer Status Register (TSR)	10-51
Figure 10-35. User SPR General 0 (USPRG0)	10-52
Figure 10-36. Fixed Point Exception Register (XER)	10-53
Figure 10-37. Zone Protection Register (ZPR)	10-54

Figure A-1. I Instruction Format	A-44
Figure A-2. B Instruction Format	A-44
Figure A-3. SC Instruction Format	A-44
Figure A-4. D Instruction Format	A-44
Figure A-5. X Instruction Format	A-45
Figure A-6. XL Instruction Format	A-45
Figure A-7. XFX Instruction Format	A-46
Figure A-8. XO Instruction Format	A-46
Figure A-9. M Instruction Format	A-46

Tables

Table 2-1. PPC405 SPRs	2-6
Table 2-2. XER[CA] Updating Instructions	2-9
Table 2-3. XER[SO,OV] Updating Instructions	2-9
Table 2-4. Time Base Registers	2-13
Table 2-5. Alignment Exception Summary	2-17
Table 2-6. Bits of the BO Field	2-25
Table 2-7. Conditional Branch BO Field	2-26
Table 2-8. Example Memory Mapping.....	2-30
Table 2-9. Privileged Instructions	2-31
Table 2-10. PPC405 Instruction Set Summary.....	2-36
Table 2-11. Implementation-specific Instructions	2-37
Table 2-12. Storage Reference Instructions	2-37
Table 2-13. Arithmetic Instructions	2-38
Table 2-14. Multiply-Accumulate and Multiply Halfword Instructions	2-39
Table 2-15. Logical Instructions	2-39
Table 2-16. Compare Instructions	2-39
Table 2-17. Branch Instructions	2-40
Table 2-18. CR Logical Instructions	2-40
Table 2-19. Rotate Instructions	2-40
Table 2-20. Shift Instructions	2-41
Table 2-21. Cache Management Instructions	2-41
Table 2-22. Interrupt Control Instructions	2-41
Table 2-23. TLB Management Instructions	2-42
Table 2-24. Processor Management Instructions	2-42
Table 3-1. MSR Contents after Reset	3-2
Table 3-2. SPR Contents After Reset	3-3
Table 4-1. Available Cache Array Sizes	4-2
Table 4-2. ICU and DCU Cache Array Organization	4-3
Table 4-3. Cache Sizes, Tag Fields, and Lines.....	4-3
Table 4-4. Priority Changes With Different Data Cache Operations	4-17
Table 5-1. Interrupt Handling Priorities	5-4
Table 5-2. Interrupt Vector Offsets	5-6
Table 5-3. ESR Alteration by Various Interrupts	5-13
Table 5-4. Register Settings during Critical Input Interrupts	5-14
Table 5-5. Register Settings during Machine Check—Instruction Interrupts	5-15
Table 5-6. Register Settings during Machine Check—Data Interrupts	5-15
Table 5-7. Register Settings during Data Storage Interrupts	5-17
Table 5-8. Register Settings during Instruction Storage Interrupts	5-18
Table 5-9. Register Settings during External Interrupts	5-19
Table 5-10. Alignment Interrupt Summary	5-19
Table 5-11. Register Settings during Alignment Interrupts	5-19
Table 5-12. ESR Usage for Program Interrupts	5-20

Table 5-13. Register Settings during Program Interrupts	5-21
Table 5-14. Register Settings during FPU Unavailable Interrupts	5-21
Table 5-15. Register Settings during System Call Interrupts	5-22
Table 5-16. Register Settings during APU Unavailable Interrupts	5-22
Table 5-17. Register Settings during Programmable Interval Timer Interrupts	5-23
Table 5-18. Register Settings during Fixed Interval Timer Interrupts	5-24
Table 5-19. Register Settings during Watchdog Timer Interrupts	5-24
Table 5-20. Register Settings during Data TLB Miss Interrupts	5-25
Table 5-21. Register Settings during Instruction TLB Miss Interrupts	5-25
Table 5-22. SRR2 during Debug Interrupts	5-26
Table 5-23. Register Settings during Debug Interrupts	5-26
Table 6-1. Time Base Access	6-3
Table 6-2. FIT Controls	6-5
Table 6-3. Watchdog Timer Controls	6-6
Table 7-1. TLB Fields Related to Page Size	7-4
Table 7-2. Protection Applied to Cache Control Instructions	7-16
Table 8-1. Debug Events	8-11
Table 8-2. DAC Applied to Cache Instructions	8-15
Table 8-3. Setting of DBSR Bits for DAC and DVC Events	8-17
Table 8-4. Comparisons Based on DBCR1[DVnM]	8-18
Table 8-5. Comparisons for Aligned DVC Accesses	8-18
Table 8-6. Comparisons for Misaligned DVC Accesses	8-19
Table 8-7. JTAG Connector Signals	8-20
Table 8-8. JTAG Instructions	8-21
Table 9-1. Implementation-Specific Instructions	9-1
Table 9-2. Operator Precedence	9-5
Table 9-3. Extended Mnemonics for addi	9-9
Table 9-4. Extended Mnemonics for addic	9-10
Table 9-5. Extended Mnemonics for addic.	9-11
Table 9-6. Extended Mnemonics for addis	9-12
Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla	9-21
Table 9-8. Extended Mnemonics for bcctr, bcctrl	9-27
Table 9-9. Extended Mnemonics for bclr, bclrl	9-30
Table 9-10. Extended Mnemonics for cmp	9-34
Table 9-11. Extended Mnemonics for cmpi	9-35
Table 9-12. Extended Mnemonics for cmpl	9-36
Table 9-13. Extended Mnemonics for cmpli	9-37
Table 9-14. Extended Mnemonics for creqv	9-41
Table 9-15. Extended Mnemonics for crnor	9-43
Table 9-16. Extended Mnemonics for cror	9-44
Table 9-17. Extended Mnemonics for crxor	9-46
Table 9-18. Transfer Bit Mnemonic Assignment.....	9-108
Table 9-19. Extended Mnemonics for mfspr	9-113
Table 9-20. Extended Mnemonics for mftb.....	9-114
Table 9-21. Extended Mnemonics for mftb	9-115
Table 9-22. Extended Mnemonics for mtrcf	9-116

Table 9-23. Extended Mnemonics for mtspr	9-120
Table 9-24. Extended Mnemonics for nor, nor.	9-139
Table 9-25. Extended Mnemonics for or, or.	9-140
Table 9-26. Extended Mnemonics for ori	9-142
Table 9-27. Extended Mnemonics for rlwimi, rlwimi.	9-146
Table 9-28. Extended Mnemonics for rlwinm, rlwinm.	9-147
Table 9-29. Extended Mnemonics for rlwnm, rlwnm.	9-150
Table 9-30. Extended Mnemonics for subf, subf., subfo, subfo.	9-176
Table 9-31. Extended Mnemonics for subfc, subfc., subfco, subfco.	9-177
Table 9-32. Extended Mnemonics for tlbre	9-185
Table 9-33. Extended Mnemonics for tlbre	9-189
Table 9-34. Extended Mnemonics for tw	9-191
Table 9-35. Extended Mnemonics for twi	9-194
Table 10-1. PPC405 General Purpose Registers.....	10-1
Table 10-2. Special Purpose Registers	10-2
Table 10-3. Time Base Registers	10-4
Table A-1. PPC405 Instruction Syntax Summary	A-1
Table A-2. PPC405 Instructions by Opcode	A-33
Table B-1. PPC405 Instruction Set Categories	B-1
Table B-2. Implementation-specific Instructions	B-1
Table B-3. Instructions in the IBM PowerPC Embedded Environment	B-5
Table B-4. Privileged Instructions	B-7
Table B-5. Extended Mnemonics for PPC405	B-10
Table B-6. Storage Reference Instructions	B-29
Table B-7. Arithmetic and Logical Instructions	B-33
Table B-8. Condition Register Logical Instructions	B-37
Table B-9. Branch Instructions	B-38
Table B-10. Comparison Instructions	B-39
Table B-11. Rotate and Shift Instructions	B-40
Table B-12. Cache Control Instructions	B-41
Table B-13. Interrupt Control Instructions	B-42
Table B-14. TLB Management Instructions	B-42
Table B-15. Processor Management Instructions	B-44
Table C-1. Cache Sizes, Tag Fields, and Lines	C-2
Table C-2. Multiply and MAC Instruction Timing	C-5
Table C-3. Instruction Cache Miss Penalties.....	C-7

About This Book

This user's manual provides the architectural overview, programming model, and detailed information about the registers, the instruction set, and operations of the IBM™ PowerPC™ 405 (PPC405 core) 32-bit RISC embedded processor core.

The PPC405 RISC embedded processor core features:

- PowerPC Architecture™
- Single-cycle execution for most instructions
- Instruction cache unit and data cache unit
- Support for little endian operation
- Interrupt interface for one critical and one non-critical interrupt signal
- JTAG interface
- Extensive development tool support

Who Should Use This Book

This book is for system hardware and software developers, and for application developers who need to understand the PPC405 core. The audience should understand embedded processor design, embedded system design, operating systems, RISC processing, and design for testability.

How to Use This Book

This book describes the PPC405 device architecture, programming model, external interfaces, internal registers, and instruction set. This book contains the following chapters, arranged in parts:

Chapter 1	Overview
Chapter 2	Programming Model
Chapter 3	Initialization
Chapter 4	Cache Operations
Chapter 5	Fixed-Point Interrupts and Exceptions
Chapter 6	Timer Facilities
Chapter 7	Memory Management
Chapter 8	Debugging
Chapter 9	Instruction Set
Chapter 10	Register Summary

This book contains the following appendixes:

Appendix A	Instruction Summary
Appendix B	Instructions by Category
Appendix C	Code Optimization and Instruction Timings

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

- Contents, on page v.
- Figures, on page xv.
- Tables, on page xviii.
- Index, on page X-1.

Conventions

The following is a list of notational conventions frequently used in this manual.

$\overline{\text{ActiveLow}}$	An overbar indicates an active-low signal.
n	A decimal number
$0xn$	A hexadecimal number
$0bn$	A binary number
$=$	Assignment
\wedge	AND logical operator
\neg	NOT logical operator
\vee	OR logical operator
\oplus	Exclusive-OR (XOR) logical operator
$+$	Twos complement addition
$-$	Twos complement subtraction, unary minus
\times	Multiplication
\div	Division yielding a quotient
$\%$	Remainder of an integer division; $(33 \% 32) = 1$.
$\ $	Concatenation
$=, \neq$	Equal, not equal relations
$<, >$	Signed comparison relations
\lessgtr, \gtrless	Unsigned comparison relations
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.
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.
leave	Leave innermost do loop or do loop specified in a leave statement.
FLD	An instruction or register field
FLD _b	A bit in a named instruction or register field
FLD _{b:b}	A range of bits in a named instruction or register field

FLD _{b,b,...}	A list of bits, by number or name, in a named instruction or register field
REG _b	A bit in a named register
REG _{b:b}	A range of bits in a named register
REG _{b,b,...}	A list of bits, by number or name, in a named register
REG[FLD]	A field in a named register
REG[FLD, FLD ...]	A list of fields in a named register
REG[FLD:FLD]	A range of fields in a named register
GPR(<i>r</i>)	General Purpose Register (GPR) <i>r</i> , where $0 \leq r \leq 31$.
(GPR(<i>r</i>))	The contents of GPR <i>r</i> , where $0 \leq r \leq 31$.
DCR(DCRN)	A Device Control Register (DCR) specified by the DCRF field in an mf dcr or mt dcr instruction
SPR(SPRN)	An SPR specified by the SPRF field in an mf spr or mt spr instruction
TBR(TBRN)	A Time Base Register (TBR) specified by the TBRF field in an mf tb instruction
GPR _s	RA, RB, ...
(Rx)	The contents of a GPR, where <i>x</i> is A, B, S, or T
(RA 0)	The contents of the register RA or 0, if the RA field is 0.
CR _{FLD}	The field in the condition register pointed to by a field of an instruction.
c _{0:3}	A 4-bit object used to store condition results in compare instructions.
ⁿ b	The bit or bit value <i>b</i> is replicated <i>n</i> times.
xx	Bit positions which are don't-cares.
CEIL(<i>x</i>)	Least integer $\geq x$.
EXTS(<i>x</i>)	The result of extending <i>x</i> on the left with sign bits.
PC	Program counter.
RESERVE	Reserve bit; indicates whether a process has reserved a block of storage.
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.
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.
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> .
EA	Effective address; the 32-bit address, derived by applying indexing or indirect addressing rules to the specified operand, that specifies a location in main storage.

EA_b	A bit in an effective address.
$EA_{b:b}$	A range of bits in an effective address.
$ROTL((RS),n)$	Rotate left; the contents of RS are shifted left the number of bits specified by n .
$MASK(MB,ME)$	Mask having 1s in positions MB through ME (wrapping if $MB > ME$) and 0s elsewhere.
$instruction(EA)$	An instruction operating on a data or instruction cache block associated with an EA.

Chapter 1. Overview

The IBM 405 32-bit reduced instruction set computer (RISC) processor core, referred to as the PPC405 core, implements the PowerPC Architecture with extensions for embedded applications.

This chapter describes:

- PPC405 core features
- The PowerPC Architecture
- The PPC405 implementation of the IBM PowerPC Embedded Environment, an extension of the PowerPC Architecture for embedded applications
- PPC405 organization, including a block diagram and descriptions of the functional units
- PPC405 registers
- PPC405 addressing modes

1.1 PPC405 Features

The PPC405 core provides high performance and low power consumption. The PPC405 RISC CPU executes at sustained speeds approaching one cycle per instruction. On-chip instruction and data caches arrays can be implemented to reduce chip count and design complexity in systems and improve system throughput.

The PowerPC RISC fixed-point CPU features:

- PowerPC User Instruction Set Architecture (UISA) and extensions for embedded applications
- Thirty-two 32-bit general purpose registers (GPRs)
- Static branch prediction
- Five-stage pipeline with single-cycle execution of most instructions, including loads/stores
- Unaligned load/store support to cache arrays, main memory, and on-chip memory (OCM)
- Hardware multiply/divide for faster integer arithmetic (4-cycle multiply, 35-cycle divide)
- Multiply-accumulate instructions
- Enhanced string and multiple-word handling
- True little endian operation
- Programmable Interval Timer (PIT), Fixed Interval Timer (FIT), and watchdog timer
- Forward and reverse trace from a trigger event
- Storage control
 - Separate, configurable, two-way set-associative instruction and data cache units; for the PPC405B3, the instruction cache array is 16KB and the data cache array is 8KB
 - Eight words (32 bytes) per cache line
 - Support for any combination of 0KB, 4KB, 8KB, and 16KB, and 32KB instruction and data cache arrays, depending on model

- Instruction cache unit (ICU) non-blocking during line fills, data cache unit (DCU) non-blocking during line fills and flushes
- Read and write line buffers
- Instruction fetch hits are supplied from line buffer
- Data load/store hits are supplied to line buffer
- Programmable ICU prefetching of next sequential line into line buffer
- Programmable ICU prefetching of non-cacheable instructions, full line (eight words) or half line (four words)
- Write-back or write-through DCU write strategies
- Programmable allocation on loads and stores
- Operand forwarding during cache line fills
- Memory Management
 - Translation of the 4GB logical address space into physical addresses
 - Independent enabling of instruction and data translation/protection
 - Page level access control using the translation mechanism
 - Software control of page replacement strategy
 - Additional control over protection using zones
 - WIU0GE (write-through, cachability, compressed/user-defined 0, guarded, endian) storage attribute control for each virtual memory region
- WIU0GE storage attribute control for thirty-two real 128MB regions in real mode
- Support for OCM that provides memory access performance identical to cache hits
- Full PowerPC floating-point unit (FPU) support using the auxiliary processor unit (APU) interface (the PPC405 does not include an FPU)
- PowerPC timer facilities
 - 64-bit time base
 - PIT, FIT, and watchdog timers
 - Synchronous external time base clock input
- Debug Support
 - Enhanced debug support with logical operators
 - Four instruction address compares (IACs)
 - Two data address compares (DACs)
 - Two data value compares (DVCs)
 - JTAG instruction to write to ICU
 - Forward or backward instruction tracing
- Minimized interrupt latency
- Advanced power management support

1.2 PowerPC Architecture

The PowerPC Architecture comprises three levels of standards:

- PowerPC User Instruction Set Architecture (UISA), including the base user-level instruction set, user-level registers, programming model, data types, and addressing modes. This is referred to as Book I of the PowerPC Architecture.
- PowerPC Virtual Environment Architecture, describing the memory model, cache model, cache-control instructions, address aliasing, and related issues. While accessible from the user level, these features are intended to be accessed from within library routines provided by the system software. This is referred to as Book II of the PowerPC Architecture.
- PowerPC Operating Environment Architecture, including the memory management model, supervisor-level registers, and the exception model. These features are not accessible from the user level. This is referred to as Book III of the PowerPC Architecture.

Book I and Book II define the instruction set and facilities available to the application programmer. Book III defines features, such as system-level instructions, that are not directly accessible by user applications. The PowerPC Architecture is described in *The PowerPC Architecture: A Specification for a New Family of RISC Processors*.

The PowerPC Architecture provides compatibility of PowerPC Book I application code across all PowerPC implementations to help maximize the portability of applications developed for PowerPC processors. This is accomplished through compliance with the first level of the architectural definition, the PowerPC UISA, which is common to all PowerPC implementations.

1.3 The PPC405 as a PowerPC Implementation

The PPC405 implements the PowerPC UISA, user-level registers, programming model, data types, addressing modes, and 32-bit fixed-point operations. The PPC405 fully complies with the PowerPC UISA. The UISA 64-bit operations are not implemented, nor are the floating point operations, unless a floating point unit (FPU) is implemented. The floating point operations, which cause exceptions, can then be emulated by software.

Most of the features of the PPC405 are compatible with the PowerPC Virtual Environment and Operating Environment Architectures, as implemented in PowerPC processors such as the 6xx/7xx family. The PPC405 also provides a number of optimizations and extensions to these layers of the PowerPC Architecture. The full architecture of the PPC405 is defined by the PowerPC Embedded Environment and the PowerPC User Instruction Set Architecture.

The primary extensions of the PowerPC Architecture defined in the Embedded Environment are:

- A simplified memory management mechanism with enhancements for embedded applications
- An enhanced, dual-level interrupt structure
- An architected DCR address space for integrated peripheral control
- The addition of several instructions to support these modified and extended resources

Finally, some of the specific implementation features of the PPC405 are beyond the scope of the PowerPC Architecture. These features are included to enhance performance, integrate functionality, and reduce system complexity in embedded control applications.

1.4 Processor Core Organization

The processor core consists of a 5-stage pipeline, separate instruction and data cache units, virtual memory management unit (MMU), three timers, debug, and interfaces to other functions.

Figure 1-1 illustrates the logical organization of the PPC405.

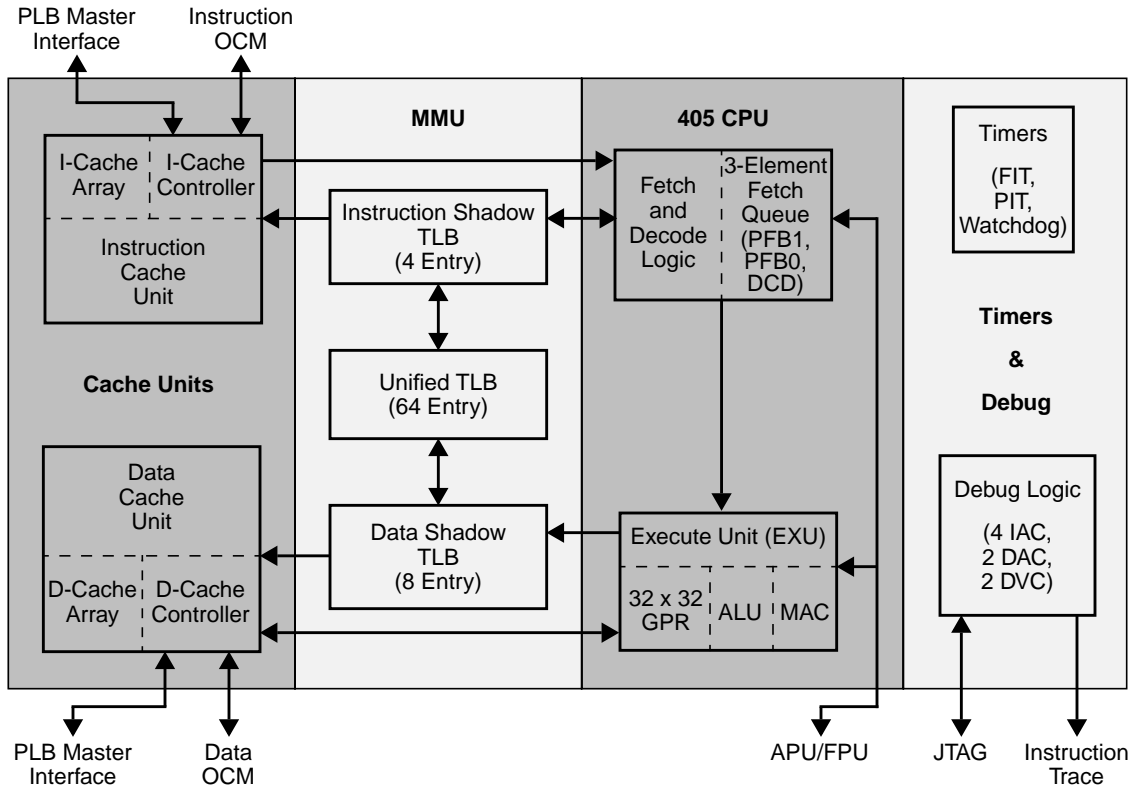


Figure 1-1. PPC405 Block Diagram

1.4.1 Instruction and Data Cache Controllers

The instruction cache unit (ICU) and data cache unit (DCU) enable concurrent accesses and minimize pipeline stalls. The storage capacity of the cache units, which can range from 0KB–32KB, depends upon the implementation. Both cache units are two-way set-associative, use a 32-byte line size. The instruction set provides a rich assortment of cache control instructions, including instructions to read tag information and data arrays. See Chapter 4, “Cache Operations,” for detailed information about the ICU and DCU.

The cache units are PLB-compliant for use in the IBM Core+ASIC program.

1.4.1.1 Instruction Cache Unit

The ICU provides one or two instructions per cycle to the execution unit (EXU) over a 64-bit bus. A line buffer (built into the output of the array for manufacturing test) enables the ICU to be accessed only once for every four instructions, to reduce power consumption by the array.

The ICU can forward any or all of the words of a line fill to the EXU to minimize pipeline stalls caused by cache misses. The ICU aborts speculative fetches abandoned by the EXU, eliminating

unnecessary line fills and enabling the ICU to handle the next EXU fetch. Aborting abandoned requests also eliminates unnecessary external bus activity to increase external bus utilization.

1.4.1.2 Data Cache Unit

The DCU transfers 1, 2, 3, 4, or 8 bytes per cycle, depending on the number of byte enables presented by the CPU. The DCU contains a single-element command and store data queue to reduce pipeline stalls; this queue enables the DCU to independently process load/store and cache control instructions. Dynamic PLB request prioritization reduces pipeline stalls even further. When the DCU is busy with a low-priority request while a subsequent storage operation requested by the CPU is stalled, the DCU automatically increases the priority of the current request to the PLB.

The DCU uses a two-line flush queue to minimize pipeline stalls caused by cache misses. Line flushes are postponed until after a line fill is completed. Registers comprise the first position of the flush queue; the line buffer built into the output of the array for manufacturing test serves as the second position of the flush queue. Pipeline stalls are further reduced by forwarding the requested word to the CPU during the line fill. Single-queued flushes are non-blocking. When a flush operation is pending, the DCU can continue to access the array to determine subsequent load or store hits. Under these conditions, load hits can occur concurrently with store hits to write-back memory without stalling the pipeline. Requests abandoned by the CPU can also be aborted by the cache controller.

Additional DCU features enable the programmer to tailor performance for a given application. The DCU can function in write-back or write-through mode, as controlled by the Data Cache Write-through Register (DCWR) or the translation look-aside buffer (TLB). DCU performance can be tuned to balance performance and memory coherency. Store-without-allocate, controlled by the SWOA field of the Core Configuration Register 0 (CCR0), can inhibit line fills caused by store misses to further reduce potential pipeline stalls and unwanted external bus traffic. Similarly, load-without-allocate, controlled by CCR0[LWOA], can inhibit line fills caused by load misses.

1.4.2 Memory Management Unit

The 4GB address space of the PPC405 is presented as a flat address space.

The MMU provides address translation, protection functions, 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 4GB logical address space into physical addresses
- Independent enabling of instruction and data translation/protection
- Page level access control using the translation mechanism
- Software control of page replacement strategy
- Additional control over protection using zones
- Storage attributes for cache policy and speculative memory access control

The MMU can be disabled under software control. If the MMU is not used, the PPC405 core provides other storage control mechanisms.

The translation lookaside buffer (TLB) is the hardware resource that controls translation and protection. It consists of 64 entries, each specifying a page to be translated. The TLB is fully

associative; a page entry can be placed anywhere in the TLB. The translation function of the MMU occurs pre-cache for data accesses. Cache tags and indexing use physical addresses for data accesses; instruction fetches are virtually indexed and physically tagged.

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 to manage TLB entries. These instructions are privileged and require the software to be executing in supervisor state. Additional TLB instructions are provided to move TLB entry fields to and from GPRs.

The MMU divides logical storage into pages. Eight page sizes (1KB, 4KB, 16KB, 64KB, 256KB, 1MB, 4MB, 16MB) are simultaneously supported, so that, at any given time, the TLB can contain entries for any combination of page sizes. For a logical to physical translation to occur, a valid entry for the page containing the logical address must be in the TLB. Addresses for which no TLB entry exists cause TLB-Miss exceptions.

To improve performance, 4 instruction-side and 8 data-side TLB entries are kept in shadow arrays. The shadow arrays prevent TLB contention. Hardware manages the replacement and invalidation of shadow-TLB entries; no system software action is required. The shadow arrays can be thought of as level 1 TLBs, with the main TLB serving as a level 2 TLB.

When address translation is enabled, the translation mechanism provides a basic level of protection. Physical addresses not mapped by a page entry are inaccessible when translation is enabled. Read access is implied by the existence of the valid entry in the TLB. The EX and WR bits in the TLB entry further define levels of access for the page, by permitting execute and write access, respectively.

The Zone Protection Register (ZPR) enables the system software to override the TLB access controls. For example, the ZPR provides a way to deny read access to application programs. The ZPR can be used to classify storage by type; access by type can be changed without manipulating individual TLB entries.

The PowerPC Architecture provides WIU0GE (write-back/write through, cachability, user-defined 0, guarded, endian) storage attributes that control memory accesses, using bits in the TLB or, when address translation is disabled, storage attribute control registers.

When address translation is enabled ($MSR[IR, DR] = 1$), storage attribute control bits in the TLB control the storage attributes associated with the current page. When address translation is disabled ($MSR[IR, DR] = 0$), bits in each storage attribute control register control the storage attributes associated with storage regions. Each storage attribute control register contains 32 fields. Each field sets the associated storage attribute for a 128MB memory region. See “Real-Mode Storage Attribute Control” on page 7-17 for more information about the storage attribute control registers.

1.4.3 Timer Facilities

The processor core contains a time base and three timers:

- Programmable Interval Timer (PIT)
- Fixed Interval Timer (FIT)
- Watchdog timer

The time base is a 64-bit counter incremented either by an internal signal equal to the CPU clock rate or by a separate external timer clock signal. No interrupts are generated when the time base rolls over.

The PIT is a 32-bit register that is decremented at the same rate as the time base is incremented. The user loads the PIT register with a value to create the desired delay. When a decrement occurs on a PIT count of 1, the timer stops decrementing, a bit is set in the Timer Status Register (TSR), and a PIT interrupt is generated. Optionally, the PIT can be programmed to reload automatically the last value written to the PIT register, after which the PIT begins decrementing again. The Timer Control Register (TCR) contains the interrupt enable for the PIT interrupt.

The FIT generates periodic interrupts based on selected bits in the time base. Users can select one of four intervals for the timer period by setting the appropriate bits in the TCR. When the selected bit in the time base changes from 0 to 1, a bit is set in the TSR and a FIT interrupt is generated. The FIT interrupt enable is contained in the TCR.

The watchdog timer generates a periodic interrupt based on selected bits in the time base. Users can select one of four time periods for the interval and the type of reset generated if the watchdog timer expires twice without an intervening clear from software.

1.4.4 Debug

The processor core 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 JTAG port can also be used for board test.

The debug modes, events, controls, and interfaces provide a powerful combination of debug facilities for hardware and software development tools.

1.4.4.1 Development Tool Support

The PPC405 supports a wide range of hardware and software development tools.

An 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 processor core.

1.4.4.2 Debug Modes

The internal, external, real-time-trace, and debug wait modes support a variety of debug tool used in embedded systems development. These debug modes are described in detail in “Debug Modes” on page 8-1.

1.4.5 Core Interfaces

The core provides a range of I/O interfaces that simplify the attachment of on-chip and off-chip devices.

1.4.5.1 Processor Local Bus

The PLB-compliant interface provides separate 32-bit address and 64-bit data buses for the instruction and data sides.

1.4.5.2 Device Control Register Bus

The Device Control Register (DCR) bus supports the attachment of on-chip registers for device control.

These registers are accessed using the **mfdcr** and **mtdcr** instructions.

1.4.5.3 Clock and Power Management

This interface supports several methods of clock distribution and power management.

1.4.5.4 JTAG

The JTAG port is enhanced to support the attachment of a debug tool such as the RISCWatch product from IBM Microelectronics. Through the JTAG test access port, a debug tool can single-step the processor and interrogate internal processor state to facilitate 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.

1.4.5.5 Interrupts

The processor core provides an interface to an on-chip interrupt controller that is logically outside the core. The interrupt controller combines asynchronous interrupt inputs from on-chip and off-chip sources and presents them to the core using a pair of interrupt signals: critical and non-critical. The sources of asynchronous interrupts are external signals, the JTAG/debug unit, and any implemented peripherals.

1.4.5.6 Auxiliary Processor Unit

The auxiliary processor unit (APU) interface supports the attachment of auxiliary processor hardware and the implementation of the associated instructions for improved performance in specialized applications.

1.4.5.7 On-Chip Memory

The on-chip memory (OCM) interface supports the implementation of instruction- and data-side memory that can be accessed at performance levels matching the cache arrays.

1.4.6 Data Types

Processor core operands are bytes, halfwords, and words. Multiple words or strings of bytes can be transferred using the load/store multiple and load/store string instructions. Data is represented in twos complement notation or in unsigned fixed-point format.

The address of a multibyte operand is always the lowest memory address occupied by that operand. Byte ordering can be selected as big endian (the lowest memory address of an operand contains its most significant byte) or as little endian (the lowest memory address of an operand contains its least

significant byte). See “Byte Ordering” on page 2-17 for more information about big and little endian operation.

1.4.7 Processor Core Register Set Summary

The processor core registers can be grouped into basic categories based on function and access mode: general purpose registers (GPRs), special purpose registers (SPRs), the machine state register (MSR), the condition register (CR), and, in Core+ASIC implementations, device control registers (DCRs).

Chapter 10, “Register Summary,” provides a register diagram and a register field description table for each register.

1.4.7.1 General Purpose Registers

The processor core contains 32 GPRs; each register contains 32 bits. The contents of the GPRs can be transferred from memory using load instructions and stored to memory using store instructions. GPRs, which are specified as operands in many instructions, can also receive instruction results and the contents of other registers.

1.4.7.2 Special Purpose Registers

Special Purpose Registers (SPRs), which are part of the PowerPC Architecture, are accessed using the **mtspr** and **mfspir** instructions. SPRs control the use of the debug facilities, timers, interrupts, storage control attributes, and other architected processor resources.

All SPRs are privileged (unavailable to user-mode programs), except the Count Register (CTR), the Link Register (LR), SPR General Purpose Registers (SPRG4–SPRG7, read-only), and the Fixed-point Exception Register (XER). Note that access to the Time Base Lower (TBL) and Time Base Upper (TBU) registers, when addressed as SPRs, is write-only and privileged. However, when addressed as Time Base Registers (TBRs), read access to these registers is not privileged. See “Time Base Registers” on page 10-4 for more information.

1.4.7.3 Machine State Register

The PPC405 contains a 32-bit Machine State Register (MSR). The contents of a GPR can be written to the MSR using the **mtmsr** instruction, and the MSR contents can be read into a GPR using the **mfmsr** instruction. The MSR contains fields that control the operation of the processor core.

1.4.7.4 Condition Register

The PPC405 contains a 32-bit Condition Register (CR). These bits are grouped into eight 4-bit fields, CR[CR0]–CR[CR7]. Instructions are provided to perform logical operations on CR fields and bits within fields and to test CR bits within fields. The CR fields, which are set by compare instructions, can be used to control branches. CR[CR0] can be set implicitly by arithmetic instructions.

1.4.7.5 Device Control Registers

DCRs, which are architecturally outside of the processor core, are accessed using the **mtdcr** and **mf dcr** instructions. DCRs are used to control, configure, and hold status for various functional units that are not part of the processor core. Although the PPC405 does not contain DCRs, the **mtdcr** and **mf dcr** instructions are provided.

The **mtdcr** and **mfdcr** instructions are privileged, for all DCRs. Therefore, all accesses to DCRs are privileged. See “Privileged Mode Operation” on page 2-30.

All DCR numbers are reserved, and should be neither read nor written, unless they are part of an IBM Core+ASIC implementation.

1.4.8 Addressing Modes

The processor core supports the following addressing modes, which enable efficient retrieval and storage of data in memory:

- Base plus displacement addressing
- Indexed addressing
- Base plus displacement addressing and indexed addressing, with update

In the base plus displacement addressing mode, an effective address (EA) is formed by adding a displacement to a base address contained in a GPR (or to an implied base of 0). The displacement is an immediate field in an instruction.

In the indexed addressing mode, the EA is formed by adding an index contained in a GPR to a base address contained in a GPR (or to an implied base of 0).

The base plus displacement and the indexed addressing modes also have a “with update” mode. In “with update” mode, the effective address calculated for the current operation is saved in the base GPR, and can be used as the base in the next operation. The “with update” mode relieves the processor from repeatedly loading a GPR with an address for each piece of data, regardless of the proximity of the data in memory.

Chapter 2. Programming Model

The programming model of the PPC405 embedded processor core describes the following features and operations:

- Memory organization and addressing, starting on page 2-1
- Registers, starting on page 2-2
- Data types and alignment, starting on page 2-16
- Byte ordering, starting on page 2-17
- Instruction processing, starting on page 2-23
- Branching control, starting on page 2-24
- Speculative accesses, starting on page 2-27
- Privileged mode operation, starting on page 2-30
- Synchronization, starting on page 2-33
- Instruction set, starting on page 2-36

2.1 User and Privileged Programming Models

The PPC405 executes programs in two modes, also referred to as states. Programs running in *privileged mode* (also referred to as the supervisor state) can access any register and execute any instruction. These instructions and registers comprise the privileged programming model. In *user mode*, certain registers and instructions are unavailable to programs. This is also called the problem state. Those registers and instructions that are available comprise the user programming model.

Privileged mode provides operating system software access to all processor resources. Because access to certain processor resources is denied in user mode, application software runs in user mode. Operating system software and other application software is protected from the effects of an errant application program.

Throughout this book, the terms user program and privileged programs are used to associate programs with one of the programming models. Registers and instructions are described as user or privileged. Privileged mode operation is described in detail in “Privileged Mode Operation” on page 2-30.

2.2 Memory Organization and Addressing

The PowerPC Architecture defines a 32-bit, 4-gigabyte (GB) flat address space for instructions and data

User’s manuals for standard products containing a PPC405 core describe the memory organizations and physical address maps of the standard products.

2.2.1 Storage Attributes

The PowerPC Architecture defines storage attributes that control data and instruction accesses. Storage attributes are provided to control cache write-through policy (the W storage attribute), cachability (the I storage attribute), memory coherency in multiprocessor environments (the M storage attribute), and guarding against speculative memory accesses (the G storage attribute). The IBM PowerPC Embedded Environment defines additional storage attributes for storage compression (the U0 storage attribute) and byte ordering (the E storage attribute).

The PPC405 core provides two control mechanisms for the W, I, U0, G, and E attributes. Because the PPC405 core does not provide hardware support for multiprocessor environments, the M storage attribute, when present, has no effect.

When the PPC405 core operates in virtual mode (address translation is enabled), each storage attribute is controlled by the W, I, U0, G, and E fields in the translation lookaside buffer (TLB) entry for each memory page. The size of memory pages, and hence the size of storage attribute control regions, is variable. Multiple sizes can be in effect simultaneously on different pages.

When the PPC405 core operates in real mode (address translation is disabled), storage attribute control registers control the corresponding storage attributes. These registers are:

- Data Cache Write-through Register (DCWR)
- Data Cache Cachability Register (DCCR)
- Instruction Cache Cachability Register (ICCR)
- Storage Guarded Register (SGR)
- Storage Little-Endian Register (SLER)
- Storage User-defined 0 Register (SU0R)

Each storage attribute control register contains 32 bits; each bit controls one of thirty-two 128MB storage attribute control regions. Bit 0 of each register controls the lowest-order region, with ascending bits controlling ascending regions in memory. The storage attributes in each storage attribute region are set independently of each other and of the storage attributes for other regions.

2.3 Registers

All PPC405 registers are listed in this section. Some of the frequently-used registers are described in detail. Other registers are covered in their respective topic chapters (for example, the cache registers are described in Chapter 4, “Cache Operations”). All registers are summarized in Chapter 10, “Register Summary.”

The registers are grouped into categories: General Purpose Registers (GPRs), Special Purpose Registers (SPRs), Time Base Registers (TBRs), the Machine State Register (MSR), the Condition Register (CR), and, in standard products, Device Control Registers (DCRs). Different instructions are used to access each category of registers.

For all registers with fields marked as *reserved*, the reserved fields should be written as 0 and read as *undefined*. That is, when writing to a register with a reserved field, write a 0 to the reserved field. When reading from a register with a reserved field, ignore that field.

Programming Note: A good coding practice is to perform the initial write to a register with reserved fields as described, 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.

Figure 2-1 on page 2-4 illustrates the registers in the user and supervisor programming models.

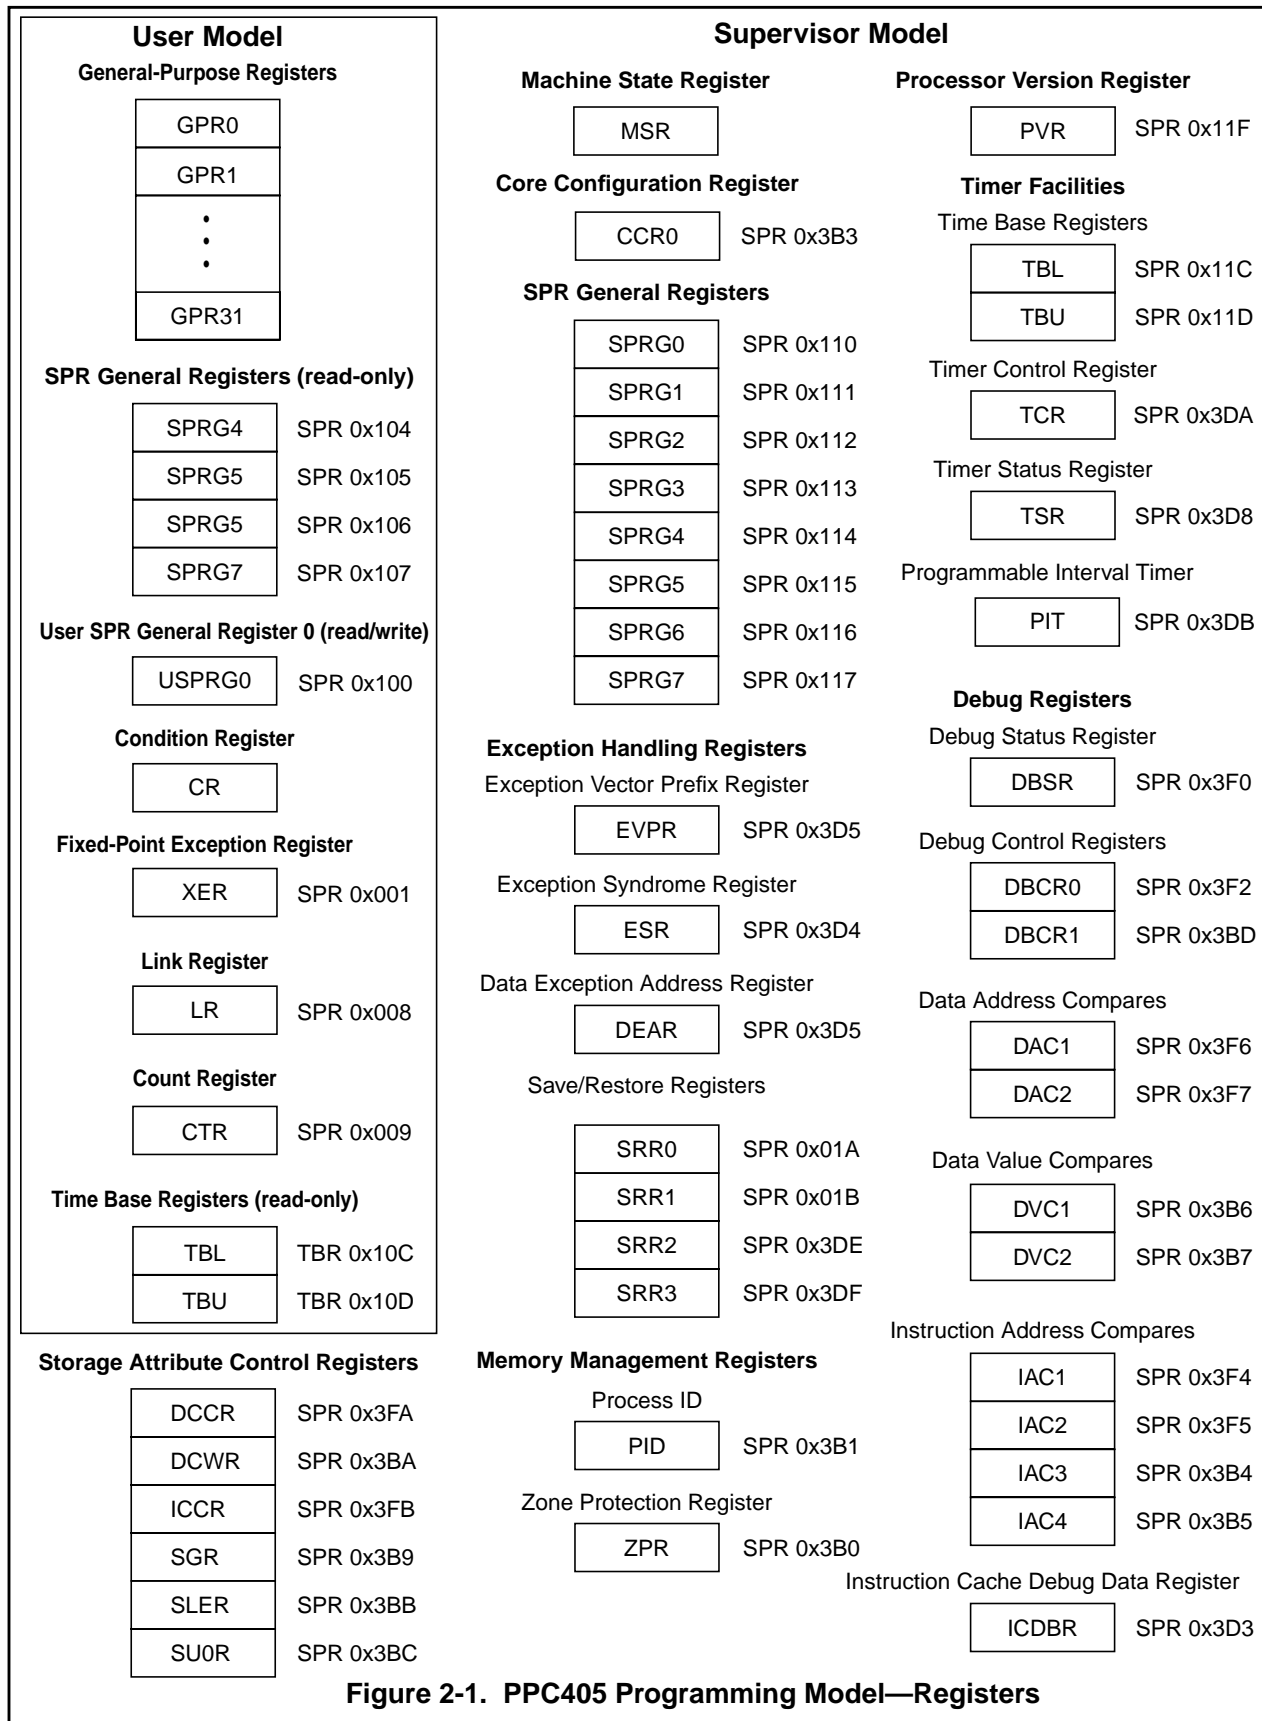


Figure 2-1. PPC405 Programming Model—Registers

2.3.1 General Purpose Registers (R0-R31)

The PPC405 core contains thirty-two 32-bit general purpose registers (GPRs). Data from memory can be read into GPRs using load instructions and the contents of GPRs can be written to memory using store instructions. Most integer instructions use GPRs for source and destination operands. See Table 10, “Register Summary,” on page 10-1 for the numbering of the GPRs.



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



2.3.2 Special Purpose Registers

Special purpose registers (SPRs), which are part of the PowerPC Architecture and the IBM PowerPC Embedded Environment, are accessed using the **mtspr** and **mfspr** instructions.

SPRs control the operation of debug facilities, timers, interrupts, storage control attributes, and other architected processor resources. Table 10, “Register Summary,” on page 10-1 shows the mnemonic, name, and number for each SPR. Table 2-1, “PPC405 SPRs,” on page 2-6 lists the PPC405 SPRs by function and indicates the pages where the SPRs are described more fully.

Except for the Link Register (LR), the Count Register (CTR), the Fixed-point Exception Register (XER), User SPR General 0 (USPRG0, and read access to SPR General 4–7 (SPRG4–SPRG7), all SPRs are privileged. As SPRs, the registers TBL and TBU are privileged write-only; as TBRs, these registers can be read in user mode. Unless used to access non-privileged SPRs, attempts to execute **mfspr** and **mtspr** instructions while in user mode cause privileged violation program interrupts. See “Privileged SPRs” on page 2-32.

Table 2-1. PPC405 SPRs

Function	Register				Access	Page
Configuration	CCR0				Privileged	4-11
Branch Control	CTR				User	2-6
	LR				User	2-7
Debug	DAC1	DAC2			Privileged	8-9
	DBCR0	DBCR1			Privileged	8-4
	DBSR				Privileged	8-7
	DVC1	DVC2			Privileged	8-10
	IAC1	IAC2	IAC3	IAC4	Privileged	8-9
	ICDBDR				Privileged	4-14
Fixed-point Exception	XER				User	2-7
General-Purpose SPR	SPRG0	SPRG1	SPRG2	SPRG3	Privileged	2-9
	SPRG4	SPRG5	SPRG6	SPRG7	User read, privileged write	2-9
	USPRG0				User	2-9
Interrupts and Exceptions	DEAR				Privileged	5-13
	ESR				Privileged	5-11
	EVPR				Privileged	5-10
	SRR0	SRR1			Privileged	5-9
	SRR2	SRR3			Privileged	5-9
Processor Version	PVR				Privileged, read-only	2-10
Storage Attribute Control	DCCR				Privileged	7-17
	DCWR				Privileged	7-17
	ICCR				Privileged	7-17
	SGR				Privileged	7-17
	SLER				Privileged	7-17
	SU0R				Privileged	7-17
Timer Facilities	TBL	TBU			Privileged, write-only	6-1
	PIT				Privileged	6-4
	TCR				Privileged	6-9
	TSR				Privileged	6-8
Zone Protection	ZPR				Privileged	7-14

2.3.2.1 Count Register (CTR)

The CTR is written from a GPR using **mtspr**. The CTR contents can be used as a loop count that is decremented and tested by some branch instructions. Alternatively, the CTR contents can specify a target address for the **bcctr** instruction, enabling branching to any address.

The CTR is in the user programming model.



Figure 2-3. Count Register (CTR)

0:31		Count	Used as count for branch conditional with decrement instructions, or as address for branch-to-counter instructions.
------	--	-------	---

2.3.2.2 Link Register (LR)

The LR is written from a GPR using **mtspr**, and by branch instructions that have the LK bit set to 1. Such branch instructions load the LR with the address of the instruction following the branch instruction. Thus, the LR contents can be used as the return address for a subroutine that was called using the branch.

The LR contents can be used as a target address for the **bclr** instruction. This allows branching to any address.

When the LR contents represent an instruction address, LR_{30:31} are assumed to be 0, because all instructions must be word-aligned. However, when LR is read using **mfspr**, all 32 bits are returned as written.

The LR is in the user programming model.



Figure 2-4. Link Register (LR)

0:31		Link Register contents	If (LR) represents an instruction address, LR _{30:31} should be 0.
------	--	------------------------	---

2.3.2.3 Fixed Point Exception Register (XER)

The XER records overflow and carry conditions generated by integer arithmetic instructions.

The Summary Overflow (SO) field is set to 1 when instructions cause the Overflow (OV) field to be set to 1. The SO field does not necessarily indicate that an overflow occurred on the most recent arithmetic operation, but that an overflow occurred since the last clearing of XER[SO]. **mtspr**(XER) sets XER[SO, OV] to the value of bit positions 0 and 1 in the source register, respectively.

Once set, XER[SO] is not reset until an **mtspr**(XER) is executed with data that explicitly puts a 0 in the SO bit, or until an **mcrxr** instruction is executed.

XER[OV] is set to indicate whether an instruction that updates XER[OV] produces a result that “overflows” the 32-bit target register. XER[OV] = 1 indicates overflow. For arithmetic operations, this occurs when an operation has a carry-in to the most-significant bit of the result that does not equal the carry-out of the most-significant bit (that is, the exclusive-or of the carry-in and the carry-out is 1).

The following instructions set XER[OV] differently. The specific behavior is indicated in the instruction descriptions in Chapter 9, “Instruction Set.”

- Move instructions:
mcrxr, mtspr(XER)
- Multiply and divide instructions:
mullwo, mullwo., divwo, divwo., divwuo, divwuo

The Carry (CA) field is set to indicate whether an instruction that updates XER[CA] produces a result that has a carry-out of the most-significant bit. XER[CA] = 1 indicates a carry.

The following instructions set XER[CA] differently. The specific behavior is indicated in the instruction descriptions in Chapter 9, “Instruction Set.”

- Move instructions
mcrxr, mtspr(XER)
- Shift-algebraic operations
sraw, srawi

The Transfer Byte Count (TBC) field is the byte count for load/store string instructions.

The XER is part of the user programming model.




Figure 2-5. Fixed Point Exception Register (XER)

0	SO	Summary Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by mtspr or by using “o” form instructions; can be <i>reset</i> by mtspr or by mcrxr .
1	OV	Overflow 0 No overflow has occurred. 0 Overflow has occurred.	Can be <i>set</i> by mtspr or by using “o” form instructions; can be <i>reset</i> by mtspr , by mcrxr , or “o” form instructions.
2	CA	Carry 0 Carry has not occurred. 1 Carry has occurred.	Can be <i>set</i> by mtspr or arithmetic instructions that update the CA field; can be <i>reset</i> by mtspr , by mcrxr , or by arithmetic instructions that update the CA field.

3:24		Reserved	
25:31	TBC	Transfer Byte Count	Used by lswx and stswx ; written by mtspr .

Table 2-2 and Table 2-3 list the PPC405 instructions that update the XER. In the tables, the syntax “[o]” indicates that the instruction has an “o” form that updates XER[SO,OV], and a “non-o” form. The syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0] (see “Condition Register (CR)” on page 2-10), and a “non-record” form.

Table 2-2. XER[CA] Updating Instructions

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

Table 2-3. 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[.]	macchwo[.] macchwso[.] macchwsuo[.] macchwuo[.] machhwo[.] machhwso[.] machhwsuo[.] machhwuo[.] maclhwo[.] maclhwsuo[.] maclhwsuo[.] maclhwo[.]	nmacchwo[.] nmacchwso[.] nmachhwo[.] nmachhwso[.] nmaclhwo[.] nmaclhwsuo[.]	mtspr mcrxr

2.3.2.4 Special Purpose Register General (SPRG0–SPRG7)

USPRG0 and SPRG0–SPRG7 are provided for general purpose software use. For example, these registers are used as temporary storage locations. For example, an interrupt handler 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 **mfspr**.

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

SPRG0–SPRG7 provide temporary storage locations. For example, an interrupt handler might save the contents of a GPR to an SPRG, and later restore the GPR from it. This is faster than performing a save/restore to memory. These registers are written by **mtspr** and read by **mf spr**.

Access to SPRG0–SPRG7 is privileged, except for read access to SPRG4–SPRG7. See “Privileged SPRs” on page 2-32 for more information.



Figure 2-6. Special Purpose Register General (SPRG0–SPRG7)

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

2.3.2.5 Processor Version Register (PVR)

The PVR is a read-only register that uniquely identifies a standard product or Core+ASIC implementation. Software can examine the PVR to recognize implementation-dependent features and determine available hardware resources.

Access to the PVR is privileged. See “Privileged SPRs” on page 2-32 for more information.




Figure 2-7. Processor Version Register (PVR)

0:11	OWN	Owner Identifier	Identifies the owner of a core
12:15	PCF	Processor Core Family	Identifies the processor core family.
16:21	CAS	Cache Array Sizes	Identifies the cache array sizes.
22:25	PCL	Processor Core Version	Identifies the core version for a specific combination of PVR[PCF] and PVR[CAS]
26:31	AID	ASIC Identifier	Assigned sequentially; identifies an ASIC function, version, and technology

2.3.3 Condition Register (CR)

The CR contains eight 4-bit fields (CR0–CR7), as shown in Figure 3-8. The fields contain conditions detected during the execution of integer or logical compare instructions, as indicated in the instruction

descriptions in Chapter 9, “Instruction Set.” The CR contents can be used in conditional branch instructions.

The CR can be modified in any of the following ways:

- **mtrf** sets specified CR fields by writing to the CR from a GPR, under control of a mask specified as an instruction field.
- **mcrf** sets a specified CR field by copying another CR field to it.
- **mcrxr** copies certain bits of the XER into a designated CR field, and then clears the corresponding XER bits.
- The “with update” forms of integer instructions implicitly update CR[CR0].
- Integer compare instructions update a specified CR field.
- Auxiliary processor instructions can update a specified CR field (including the implicit update of CR[CR1] by certain floating-point operations).
- The CR-logical instructions update a specified CR bit with the result of a logical operation on a specified pair of CR bit fields.
- Conditional branch instructions can test a CR bit as one of the branch conditions.

If a CR field is set by a compare instruction, the bits are set as described in “CR Fields after Compare Instructions.”

The CR is part of the user programming model.




Figure 2-8. 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

2.3.3.1 CR Fields after Compare Instructions

Compare instructions compare the values of two registers. The two types of compare instructions, *arithmetic* and *logical*, are distinguished by the interpretation given to the 32-bit values. For *arithmetic*

compares, the values are considered to be signed, where 31 bits represent the magnitude and the most-significant bit is a sign bit. For *logical* compares, the values are considered to be unsigned, so all 32 bits represent magnitude. There is no sign bit. As an example, consider the comparison of 0 with 0xFFFFFFFF. In an *arithmetic* compare, 0 is larger, because 0xFFFF FFFF represents -1 ; in a *logical* compare, 0xFFFFFFFF is larger.

A compare instruction can direct its CR update to any CR field. The first data operand of a compare instruction specifies a GPR. The second data operand specifies another GPR, or immediate data derived from the IM field of the immediate instruction form. The contents of the GPR specified by the first data operand are compared with the contents of the GPR specified by the second data operand (or with the immediate data). See descriptions of the compare instructions (page 9-34 through page 9-37) for precise details.

After a compare, the specified CR field is interpreted as follows:

LT (bit 0)	The first operand is less than the second operand.
GT (bit 1)	The first operand is greater than the second operand.
EQ (bit 2)	The first operand is equal to the second operand.
SO (bit 3)	Summary overflow; a copy of XER[SO].

2.3.3.2 The CR0 Field

After the execution of compare instructions that update CR[CR0], CR[CR0] is interpreted as described in “CR Fields after Compare Instructions” on page 2-11. The “dot” forms of arithmetic and logical instructions also alter CR[CR0]. After most instructions that update CR[CR0], the bits of CR0 are interpreted as follows:

LT (bit 0)	Less than 0; set if the most-significant bit of the 32-bit result is 1.
GT (bit 1)	Greater than 0; set if the 32-bit result is non-zero and the most-significant bit of the result is 0.
EQ (bit 2)	Equal to 0; set if the 32-bit result is 0.
SO (bit 3)	Summary overflow; a copy of XER[SO] at instruction completion.

The CR[CR0]_{LT,GT,EQ} subfields are set as the result of an algebraic comparison of the instruction result to 0, regardless of the type of instruction that sets CR[CR0]. If the instruction result is 0, the EQ subfield is set to 1. If the result is not 0, either LT or GT is set, depending on the value of the most-significant bit of the result.

When updating CR[CR0], the most significant bit of an instruction result is considered a sign bit, even for instructions that produce results that are not usually thought of as signed. For example, logical instructions such as **and.**, **or.**, and **nor.** update CR[CR0]_{LT,GT,EQ} using such an arithmetic comparison to 0, although the result of such a logical operation is not actually an arithmetic result.

If an arithmetic overflow occurs, the “sign” of an instruction result indicated in CR[CR0]_{LT,GT,EQ} 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 two's-complement number in a 32-bit register, an overflow occurs and CR[CR0]_{LT,SO} are set, although the infinitely precise result of the add is positive.

Adding the largest 32-bit twos-complement negative number, 0x8000 0000, to itself results in an arithmetic overflow and 0x0000 0000 is recorded in the target register. CR[CR0]_{EQ, SO} is set, indicating a result of 0, but the infinitely precise result is negative.

The CR[CR0]_{SO} subfield is a copy of XER[SO]. Instructions that do not alter the XER[SO] bit cannot cause an overflow, but even for these instructions CR[CR0]_{SO} is a copy of XER[SO].

Some instructions set CR[CR0] differently or do not specifically set any of the subfields. These instructions include:

- Compare instructions
cmp, cmpi, cmpl, cmpli
- CR logical instructions
crand, crandc, creqv, crnand, crnor, cror, crorc, crxor, mcrf
- Move CR instructions
mtrcf, mcrxr
- **stwcx.**

The instruction descriptions provide detailed information about how the listed instructions alter CR[CR0].

2.3.4 The Time Base

The PowerPC Architecture provides a 64-bit time base. “Time Base” on page 6-1 describes the architected time base. Access to the time base is through two 32-bit time base registers (TBRs). The least-significant 32 bits of the time base are read from the Time Base Lower (TBL) register and the most-significant 32 bits are read from the Time Base Upper (TBU) register.

User-mode access to the time base is read-only, and there is no explicitly privileged read access to the time base.

The **mftb** instruction reads from TBL and TBU. Writing the time base is accomplished by moving the contents of a GPR to a pair of SPRs, which are also called TBL and TBU, using **mtspr**.

Table 2-4 shows the mnemonics and names of the TBRs.

Table 2-4. Time Base Registers

Mnemonic	Register Name	Access
TBL	Time Base Lower (Read-only)	Read-only
TBU	Time Base Upper (Read-only)	Read-only

2.3.5 Machine State Register (MSR)

The Machine State Register (MSR) controls processor core functions, such as the enabling or disabling of interrupts and address translation.

The MSR is written from a GPR using the **mtmsr** instruction. The contents of the MSR can be read into a GPR using the **mfmsr** instruction. MSR[EE] is set or cleared using the **wrtee** or **wrteei** instructions.

The MSR contents are automatically saved, altered, and restored by the interrupt-handling mechanism. See “Machine State Register (MSR)” on page 5-7.




Figure 2-9. Machine State Register (MSR)

0:5		Reserved	
6	AP	Auxiliary Processor Available 0 APU not available. 1 APU available.	
7:11		Reserved	
12	APE	APU Exception Enable 0 APU exception disabled. 1 APU exception enabled.	
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 interrupts are disabled. 1 Critical interrupts are enabled.	Controls the critical interrupt input and watchdog timer first time-out interrupts.
15		Reserved	
16	EE	External Interrupt Enable 0 Asynchronous interrupts are disabled. 1 Asynchronous interrupts are enabled.	Controls the non-critical external interrupt input, PIT, and FIT interrupts.
17	PR	Problem State 0 Supervisor state (all instructions allowed). 1 Problem state (some instructions not allowed).	
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 Debug wait mode is disabled. 1 Debug wait mode is enabled.
22	DE	Debug Interrupts 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	IR	Instruction Relocate 0 Instruction address translation is disabled. 1 Instruction address translation is enabled.
27	DR	Data Relocate 0 Data address translation is disabled. 1 Data address translation is enabled.
28:31		Reserved

2.3.6 Device Control Registers

Device Control Registers (DCRs), on-chip registers that exist architecturally outside the processor core, are not part of the IBM PowerPC Embedded Environment. The Embedded Environment simply defines the existence of a DCR address space and the instructions that access the DCRs, but does not define any DCRs. The instructions that access the DCRs are **mtdcr** (move to device control register) and **mfdcr** (move from device control register).

DCRs are used to control the operations of on-chip buses, peripherals, and some processor behavior.

2.4 Data Types and Alignment

The data types consist of bytes (eight bits), halfwords (two bytes), words (four bytes), and strings (1 to 128 bytes). Figure 2-10 shows the byte, halfword, and word data types and their bit and byte definitions for big endian representations of values. Note that PowerPC bit numbering is reversed from industry conventions; bit 0 represents the most significant bit of a value.




Figure 2-10. PPC405 Data Types

Data is represented in either twos-complement notation or in an unsigned integer format; data representation is independent of alignment issues.

The address of a data object is always the lowest address of any byte comprising the object.

All instructions are words, and are word-aligned (the lowest byte address is divisible by 4).

2.4.1 Alignment for Storage Reference and Cache Control Instructions

The storage reference instructions (loads and stores; see Table 2-12, “Storage Reference Instructions,” on page 2-37) move data to and from storage. The data cache control instructions listed in Table 2-21, “Cache Management Instructions,” on page 2-41, control the contents and operation of the data cache unit (DCU). Both types of instructions form an effective address (EA). The method of calculating the EA for the storage reference and cache control instructions is detailed in the description of those instructions. See Chapter 9, “Instruction Set,” for more information.

Cache control instructions ignore the five least significant bits of the EA; no alignment restrictions exist in the DCU because of EAs. However, storage control attributes can cause alignment exceptions. When data address translation is disabled and a **dcbz** instruction references a storage region that is non-cachable, or for which write-through caching is the write strategy, an alignment exception is taken. Such exceptions result from the storage control attributes, not from EA alignment. The alignment exception enables system software to emulate the write-through function.

Alignment requirements for the storage reference instructions and the **dcread** instruction depend on the particular instruction. Table 2-5, “Alignment Exception Summary,” on page 2-17, summarizes the instructions that cause alignment exceptions.

The data targets of instructions are of types that depend upon the instruction. The load/store instructions have the following “natural” alignments:

- Load/store word instructions have word targets, word-aligned.
- Load/ store halfword instructions have halfword targets, halfword-aligned.
- Load/store byte instructions have byte targets, byte-aligned (that is, any alignment).

Misalignments are addresses that are not naturally aligned on data type boundaries. An address not divisible by four is misaligned with respect to word instructions. An address not divisible by two is misaligned with respect to halfword instructions. The PPC405 core implementation handles misalignments within and across word boundaries, but there is a performance penalty because additional cycles are required.

2.4.2 Alignment and Endian Operation

The endian storage control attribute does not affect alignment behavior. In little endian storage regions, the alignment of data is treated as it is in big endian storage regions; no special alignment exceptions occur when accessing data in little endian storage regions. Note that the alignment exceptions that apply to big endian region accesses also apply to little endian storage region accesses.

2.4.3 Summary of Instructions Causing Alignment Exceptions

Table 2-5 summarizes the instructions that cause alignment exceptions and the conditions under which the alignment exceptions occur.

Table 2-5. Alignment Exception Summary

Instructions Causing Alignment Exceptions	Conditions
dcbz	EA in non-cachable or write-through storage
dcread, lwarx, stwcx.	EA not word-aligned
APU load/store halfword	EA not halfword-aligned
APU load/store word	EA not word-aligned
APU load/store doubleword	EA not word-aligned

2.5 Byte Ordering

The following discussion describes the “endianness” of the PPC405, which, by default and in normal use is “big endian.”

If scalars (individual data items and instructions) were indivisible, “byte ordering” would not be a concern. It is meaningless to consider the order of bits or groups of bits within a byte, the smallest addressable unit of storage; 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 byte order arise.

For a machine in which the smallest addressable unit of storage is the 32-bit word, there is no question of the ordering of bytes within words. All transfers of individual scalars between registers and storage are of words, and the address of the byte containing the high-order eight bits of a scalar is the same as the address of any other byte of the scalar.

For the PowerPC Architecture, as for most computer architectures currently implemented, the smallest addressable unit of storage is the 8-bit byte. Other 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 is stored in four consecutive byte addresses. It thus becomes meaningful to discuss the order of the byte addresses with respect to the value of the scalar: that is, 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 commonly used:

- The ordering that assigns the lowest address to the highest-order (“leftmost”) 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” of the scalar, considered as a binary number, comes first in storage.

- The ordering that assigns the lowest address to the lowest-order (“rightmost”) 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” of the scalar, considered as a binary number, comes first in storage.

2.5.1 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 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.

2.5.1.1 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).

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

2.5.1.2 Little Endian Mapping

Structure *s* is shown mapped little endian.

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

2.5.2 Support for Little Endian Byte Ordering

This book describes the processor as if it operated only in a big endian fashion. In fact, the IBM PowerPC Embedded Environment also supports little endian operation.

The PowerPC little endian mode, defined in the PowerPC Architecture, is not implemented.

2.5.3 Endian (E) Storage Attribute

The endian (E) storage attribute supports direct connection of the PPC405 core to little endian peripherals and to memory containing little endian instructions and data. For every storage reference (instruction fetch or load/store access), an E storage attribute is associated with the storage region of the reference. The E attribute specifies whether that region is organized as big endian (E = 0) or little endian (E = 1).

When address translation is enabled ($MSR[IR] = 1$ or $MSR[DR] = 1$), the E field in the corresponding TLB entry controls the endianness of a memory region. When address translation is disabled ($MSR[IR] = 0$ or $MSR[DR] = 0$), the SLER controls the endianness of a memory region.

Bytes in storage that are accessed as little endian are arranged in true little endian format. The PPC405 does not support the little endian mode defined in the PowerPC architecture and used in PPC401xx and PPC403xx processors. Furthermore, no address modification is performed when accessing storage regions programmed as little endian. Instead, the PPC405 reorders the bytes as they are transferred between the processor and memory.

The on-the-fly reversal of bytes in little endian storage regions is handled in one of two ways, depending on whether the storage access is an instruction fetch or a data access (load/store). The following sections describe byte reordering for the two kinds of storage accesses.

2.5.3.1 Fetching Instructions from Little Endian Storage Regions

Instructions are words (four bytes) that are aligned on word boundaries in memory. As such, instructions in a big endian memory region are arranged with the most significant byte (MSB) of the instruction word at the lowest address.

Consider the big endian mapping of instruction p at address 00, where, for example, $p = \text{add } r7, r7, r4$:

MSB			LSB
0x00	0x01	0x02	0x03

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

LSB			MSB
0x00	0x01	0x02	0x03

When an instruction is fetched from memory, the instruction must be placed in the instruction queue in the proper order. The execution unit assumes that the MSB of an instruction word is at the lowest address. Therefore, when instructions are fetched from little endian storage regions, the four bytes of an instruction word are reversed before the instruction is decoded. In the PPC405 core, the byte reversal occurs between memory and the instruction cache unit (ICU). The ICU always stores instructions in big endian format, regardless of whether the memory region containing the instruction is programmed as big endian or little endian. Thus, the bytes are already in the proper order when an instruction is transferred from the ICU to the decode stage of the pipeline.

If a storage region is reprogrammed from one endian format to the other, the storage region must be reloaded with program and data structures in the appropriate endian format. If the endian format of instruction memory changes, the ICU must be made coherent with the updates. The ICU must be invalidated and the updated instruction memory using the new endian format must be fetched so that the proper byte ordering occurs before the new instructions are placed in the ICU.

2.5.3.2 Accessing Data in Little Endian Storage Regions

Unlike instruction fetches from little endian storage regions, data accesses from little endian storage regions are *not* byte-reversed between memory and the DCU. Data byte ordering, in memory, depends on the data type (byte, halfword, or word) of a specific data item. It is only when moving a data item *of a specific type* from or to a GPR that it becomes known what type of byte reversal is required. Therefore, byte reversal during load/store accesses is performed between the DCU and the GPR.

When accessing data in a little endian storage region:

- For byte loads/stores, no reordering occurs.
- For halfword loads/stores, bytes are reversed within the halfword.
- For word loads/stores, bytes are reversed within the word.

Note that this applies, regardless of data alignment.

The big endian and little endian mappings of the structure *s*, shown in “Structure Mapping Examples” on page 2-18, demonstrate how the size of an item determines its byte ordering. For example:

- 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.

In little endian storage regions, the alignment of data is treated as it is in big endian storage regions. Unlike PowerPC little endian mode, no special alignment exceptions occur when accessing data in little endian storage regions.

2.5.3.3 PowerPC Byte-Reverse Instructions

For big endian storage regions, normal load/store instructions move the more significant bytes of a register to and from the lower-numbered memory addresses. The load/store with byte-reverse instructions move the more significant bytes of the register to and from the higher numbered memory addresses.

As Figure 2-11 through Figure 2-14 illustrate, a normal store to a big endian storage region is the same as a byte-reverse store to a little endian storage region. Conversely, a normal store to a little endian storage region is the same as a byte-reverse store to a big endian storage region.

Figure 2-11 illustrates the contents of a GPR and memory (starting at address 00) after a normal load/store in a big endian storage region.




Figure 2-11. Normal Word Load or Store (Big Endian Storage Region)

Note that the results are identical to the results of a load/store with byte-reverse in a little endian storage region, as illustrated in Figure 2-12.




Figure 2-12. Byte-Reverse Word Load or Store (Little Endian Storage Region)

Figure 2-13 illustrates the contents of a GPR and memory (starting at address 00) after a load/store with byte-reverse in a big endian storage region.




Figure 2-13. Byte-Reverse Word Load or Store (Big Endian Storage Region)

Note that the results are identical to the results of a normal load/store in a little endian storage region, as illustrated in Figure 2-14.




Figure 2-14. Normal Word Load or Store (Little Endian Storage Region)

The E storage attribute augments the byte-reverse load/store instructions in two important ways:

- The load/store with byte-reverse instructions do not solve the problem of fetching instructions from a storage region in little endian format.
Only the endian storage attribute mechanism supports the fetching of little endian program images.
- Typical compilers cannot make general use of the byte-reverse load/store instructions, so these instructions are ordinarily used only in device drivers written in hand-coded assembler.

Compilers can, however, take full advantage of the endian storage attribute mechanism, enabling application programmers working in a high-level language, such as C, to compile programs and data structures into little endian format.

2.6 Instruction Processing

The instruction pipeline, illustrated in Figure 2-15, contains three queue locations: prefetch buffer 1 (PFB1), prefetch buffer 0 (PFB0), and decode (DCD). This queue implements a pipeline with the following functional stages: fetch, decode, execute, write-back and load write-back. Instructions are fetched from the instruction cache unit (ICU), placed in the instruction queue, and eventually dispatched to the execution unit (EXU).

Instructions are fetched from the ICU at the request of the EXU. Cachable instructions are forwarded directly to the instruction queue and stored in the ICU cache array. Non-cachable instructions are also forwarded directly to the instruction queue, but are not stored in the ICU cache array. Fetched instructions drop to the empty queue location closest to the EXU. When there is room in the queue, instructions can be returned from the ICU two at a time. If the queue is empty and the ICU is returning two instructions, one instruction drops into DCD while the other drops into PFB0. PFB1 buffers instructions when the pipeline stalls.

Branch instructions are examined in DCD and PFB0 while all other instructions are decoded in DCD. All instructions must pass through DCD before entering the EXU. The EXU contains the execute, write-back and load write-back stages of the pipe. The results of most instructions are calculated during the execute stage and written to the GPR file during the write back stage. Load instructions write the GPR file during the load write-back stage.




Figure 2-15. PPC405 Instruction Pipeline

2.7 Branch Processing

The PPC405, which provides a variety of conditional and unconditional branching instructions, uses the branch prediction techniques described in “Branch Prediction” on page 3-35.

2.7.1 Unconditional Branch Target Addressing Options

The unconditional branches (**b**, **ba**, **bl**, **bla**) carry the displacement to the branch target address as a signed 26-bit value (the 24-bit LI field right-extended with 0b00). The displacement enables unconditional branches to cover an address range of $\pm 32\text{MB}$.

For the relative (AA = 0) forms (**b**, **bl**), the target address is the current instruction address (CIA, the address of the branch instruction) plus the signed displacement.

For the absolute (AA = 1) forms (**ba**, **bla**), the target address is 0 plus the signed displacement. If the sign bit (LI[0]) is 0, the displacement is the target address. If the sign bit is 1, the displacement is a negative value and wraps to the highest memory addresses. For example, if the displacement is 0x3FF FFFC (the 26-bit representation of -4), the target address is 0xFFFF FFFC (0 – 4B, or 4 bytes below the top of memory).

2.7.2 Conditional Branch Target Addressing Options

The conditional branches (**bc**, **bca**, **bcl**, **bcla**) carry the displacement to the branch target address as a signed 16-bit value (the 14-bit BD field right-extended with 0b00). The displacement enables conditional branches to cover an address range of $\pm 32\text{KB}$.

For the relative (AA = 0) forms (**bc**, **bcl**), the target address is the CIA plus the signed displacement.

For the absolute (AA = 1) forms (**bca**, **bcla**), the target address is 0 plus the signed displacement. If the sign bit (BD[0]) is 0, the displacement is the target address. If the sign bit is 1, the displacement is negative and wraps to the highest memory addresses. For example, if the displacement is 0xFFFFC (the 16-bit representation of -4), the target address is 0xFFFF FFFC (0 - 4B, or 4 bytes from the top of memory).

2.7.3 Conditional Branch Condition Register Testing

Conditional branch instructions can test a CR bit. The value of the BI field specifies the bit to be tested (bit 0–31). The BO field controls whether the CR bit is tested, as described in the following section.

2.7.4 BO Field on Conditional Branches

The BO field of the conditional branch instruction specifies the conditions used to control branching, and specifies how the branch affects the CTR.

Conditional branch instructions can 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 this option is selected, the condition is satisfied (branch can occur) if CR[BI] = BO[1].

Conditional branch instructions can decrement the CTR by one, and after the decrement, test the CTR value. This option is selected when BO[2] = 0. If this option is selected, BO[3] specifies the condition that must be satisfied to allow a branch to be taken. If BO[3] = 0, CTR ≠ 0 is required for a branch to occur. If BO[3] = 1, CTR = 0 is required for a branch to occur.

If BO[2] = 1, the contents of the CTR are left unchanged, and the CTR does not participate in the branch condition test.

Table 2-6 summarizes the usage of the bits of the BO field. BO[4] is further discussed in “Branch Prediction.”

Table 2-6. Bits of the BO Field

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 Test for CR[BI] = 0. 1 Test for CR[BI] = 1.
BO[2]	CTR Test Control 0 Decrement CTR by one and test whether CTR satisfies the condition specified by BO[3]. 1 Do not change CTR, do not test CTR.
BO[3]	CTR Test Value 0 Test for CTR ≠ 0. 1 Test for CTR = 0.
BO[4]	Branch Prediction Reversal 0 Apply standard branch prediction. 1 Reverse the standard branch prediction.

Table 2-7 lists specific BO field contents, and the resulting actions; z represents a mandatory value of 0, and y is a branch prediction option discussed in “Branch Prediction.”

Table 2-7. Conditional Branch BO Field

BO Value	Description
0000y	Decrement the CTR, then branch if the decremented CTR $\neq 0$ 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 $\neq 0$ 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 $\neq 0$.
1z01y	Decrement the CTR, then branch if the decremented CTR = 0.
1z1zz	Branch always.

2.7.5 Branch Prediction

Conditional branches present a problem to the instruction fetcher. A branch might be taken. The branch EXU attempts to predict whether or not a branch is taken before all information necessary to determine the branch direction is available. This decision is called a *branch prediction*. The fetcher can then prefetch instructions starting at the predicted branch target address. If the prediction is correct, time is saved because the branched-to instruction is available in the instruction queue. Otherwise, the instruction pipeline stalls while the correct instruction is fetched into the instruction queue. To be effective, branch prediction must be correct most of the time.

The PowerPC Architecture enables software to reverse the default branch prediction, which is defined 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 the sign bit of the displacement for conditional branch (**bc**) instructions, and 0 for **bclr** and **bcctr** instructions.

$(\text{BO}[0] \wedge \text{BO}[2]) = 1$ only when the conditional branch tests nothing (the “branch always” condition). Obviously, the branch should be predicted taken for this case.

If the branch tests anything, $(\text{BO}[0] \wedge \text{BO}[2]) = 0$, and s entirely controls the prediction. The default prediction for this case was decided by considering the relative form of **bc**, which is commonly used at the end of loops to control the number of times that a loop is executed. The branch is taken every 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 $s = 1$.

If branch displacements are positive ($s = 0$), the branch is predicted not taken. If the branch instruction is any form of **bclr** or **bcctr** except the “branch always” forms, 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 “Unconditional Branch Target Addressing Options” on page 2-24, if the algebraic sign of the displacement is negative ($s = 1$), the branch target address is in high memory. If

the algebraic sign of the displacement is positive ($s = 0$), the branch target address 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 default prediction is taken, and for the low memory case the default prediction is not taken.

BO[4] is the *prediction reversal bit*. If BO[4] = 0, the default prediction is applied. If BO[4] = 1, the reverse of the standard prediction is applied. For the cases in Table 3-17 where BO[4] = y , software can reverse the default prediction. 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.

The PowerPC Architecture requires assemblers to provide a way to conveniently control branch prediction. For any conditional branch mnemonic, a suffix may be added to the mnemonic to control prediction, as follows:

- + Predict branch to be taken
- Predict branch to be not taken

For example, **bcctr+** causes BO[4] to be set appropriately to force the branch to be predicted taken.

2.8 Speculative Accesses

The PowerPC 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 which is not required by a sequential execution model.

For example, prefetching instructions beyond an undetermined conditional branch is a speculative fetch; if the branch is not in the predicted direction, the program, as executed, never needs the instructions from the predicted path.

Sometimes speculative accesses are inappropriate. For example, attempting to fetch instructions from addresses that cannot contain instructions can cause problems. To protect against errant accesses to “sensitive” memory or I/O devices, the PowerPC Architecture provides the G (guarded) storage attribute, which can be used to specify memory pages from which speculative accesses are prohibited. (Actually, speculative accesses to guarded storage are allowed in certain limited circumstances; if an instruction in a cache block will be executed, the rest of the cache block can be speculatively accessed.)

2.8.1 Speculative Accesses in the PPC405

The PPC405 does not perform speculative loads.

Two methods control speculative instruction fetching. If instruction address translation is enabled (MSR[IR] = 1), the G (guarded) field in the translation lookaside buffer (TLB) entries controls speculative accesses.

If instruction address translation is disabled (MSR[IR] = 0), the Storage Guarded Register (SGR) controls speculative accesses for regions of memory. When a region is guarded (speculative fetching is disallowed), instruction prefetching is disabled for that region. A fetch request must be completely resolved (no longer speculative) before it is issued. There is a considerable performance penalty for fetching from guarded storage, so guarding should be used only when required.

Note that, following any reset, the PPC405 core operates with all of storage guarded.

Note that when address translation is enabled, attempts to fetch from guarded storage result in instruction storage exceptions. Guarded memory is in most often needed with peripheral status registers that are cleared automatically after being read, because an unintended access resulting from a speculative fetch would cause the loss of status information. Because the MMU provides 64 pages with a wide range of page sizes as small as 1KB, fetching instructions from guarded storage should be unnecessary.

2.8.1.1 Prefetch Distance Down an Unresolved Branch Path

The fetcher will speculatively access up to 19 instructions down a predicted branch path, whether taken or sequential, regardless of cachability.

2.8.1.2 Prefetch of Branches to the CTR and Branches to the LR

When the instruction fetcher predicts that a **bctr** or **blr** instruction will be taken, the fetcher does not attempt to fetch an instruction from the target address in the CTR or LR if an executing instruction updates the register ahead of the branch. (See “Instruction Processing” on page 2-23 for a description of the instruction pipeline). The fetcher recognizes that the CTR or LR contains data left from an earlier use and that such data is probably not valid.

In such cases, the fetcher does not fetch the instruction at the target address until the instruction that is updating the CTR or LR completes. Only then are the “correct” CTR or LR contents known. This prevents the fetcher from speculatively accessing a completely “random” address. After the CTR or LR contents are known to be correct, the fetcher accesses no more than five instructions down the sequential or taken path of an unresolved branch, or at the address contained in the CTR or LR.

2.8.2 Preventing Inappropriate Speculative Accesses

A memory-mapped I/O peripheral, such as a serial port having a status register that is automatically reset when read provides a simple example of storage that should not be speculatively accessed. If code is in memory at an address adjacent to the peripheral (for example, code goes from 0x0000 0000 to 0x0000 0FFF, and the peripheral is at 0x0000 1000), prefetching past the end of the code will read the peripheral.

Guarding storage also prevents prefetching past the end of memory. If the highest memory address is left unguarded, the fetcher could attempt to fetch past the last valid address, potentially causing machine checks on the fetches from invalid addresses. While the machine checks do not actually cause an exception until the processor attempts to execute an instruction at an invalid address, some systems could suffer from the attempt to access such an invalid address. For example, an external memory controller might log an error.

System designers can avoid problems from speculative fetching without using the guarded storage attributes. The rest of this section describes ways to prevent speculative instruction fetches to sensitive addresses in unguarded memory regions.

2.8.2.1 Fetching Past an Interrupt-Causing or Interrupt-Returning Instruction

Suppose a **bctr** or **blr** instruction closely follows an interrupt-causing or interrupt-returning instruction (**sc**, **rfi**, or **rftci**). The fetcher does not prevent speculatively fetching past one of these instructions. In other words, the fetcher does not treat the interrupt-causing and interrupt-returning instructions specially when deciding whether to predict down a branch path. Instructions after an **rftci**, for example, are considered to be on the determined branch path.

To understand the implications of this situation, consider the code sequence:

```
handler:  aaa
          bbb
          rfi
subroutine: bctr
```

When executing the interrupt handler, the fetcher does not recognize the **rfi** as a break in the program flow, and speculatively fetches the target of the **bctr**, which is really the first instruction of a subroutine that has not been called. Therefore, the CTR might contain an invalid pointer.

To protect against such a prefetch, the software must insert an unconditional branch hang (**b \$**) just after the **rfi**. This prevents the hardware from prefetching the invalid target address used by **bctr**.

Consider also the above code sequence, with the **rfi** instruction replaced by an **sc** instruction used to initialize the CTR with the appropriate value for the **bctr** to branch to, upon return from the system call. The **sc** handler returns to the instruction following the **sc**, which can't be a branch hang. Instead, software could put a **mtctr** just before the **sc** to load a non-sensitive address into the CTR. This address will be used as the prediction address before the **sc** executes. An alternative would be to put a **mfctr** or **mtctr** between the **sc** and the **bctr**; the **mtctr** prevents the fetcher from speculatively accessing the address contained in the CTR before initialization.

2.8.2.2 Fetching Past **tw** or **twi** Instructions

The interrupt-causing instructions, **tw** and **twi**, do not require the special handling described in “Fetching Past an Interrupt-Causing or Interrupt-Returning Instruction” on page 2-28. These instructions are typically used by debuggers, which implement software breakpoints by substituting a trap instruction for the instruction originally at the breakpoint address. In a code sequence **mtlr** followed by **blr** (or **mtctr** followed by **bctr**), replacement of **mtlr/mtctr** by **tw** or **twi** leaves the LR/CTR uninitialized. It would be inappropriate to fetch from the **blr/bctr** target address. This situation is common, and the fetcher is designed to prevent the problem.

2.8.2.3 Fetching Past an Unconditional Branch

When an unconditional branch is in DCD in the instruction queue, the fetcher recognizes that the sequential instructions following the branch are unnecessary. These sequential addresses are not accessed. Addresses at the branch target are accessed instead.

Therefore, placing an unconditional branch just before the start of a sensitive address space (for example, at the “end” of a memory area that borders an I/O device) guarantees that addresses in the sensitive area will not be speculatively fetched.

2.8.2.4 Suggested Locations of Memory-Mapped Hardware

The preferred method of protecting memory-mapped hardware from inadvertent access is to use address translation, with hardware isolated to guarded pages (the G storage attribute in the associated TLB entry is set to 1.) The pages can be as small as 1KB. Code should never be stored in such pages.

If address translation is disabled, the preferred protection method is to isolate memory-mapped hardware into regions guarded using the SGR. Code should never be stored in such regions. The disadvantage of this method, compared to the preferred method, is that each region guarded by the SGR consumes 128MB of the address space.

Table 2-8 shows two address regions of the PPC405 core. Suppose a system designer can map all I/O devices and all ROM and SRAM devices into any location in either region. The choices made by the designer can prevent speculative accesses to the memory-mapped I/O devices.

Table 2-8. Example Memory Mapping

0x7800 0000 – 0x7FFF FFFF (SGR bit 15)	128MB Region 2
0x7000 0000 – 0x77FF FFFF (SGR bit 14)	128MB Region 1

A simple way to avoid the problem of speculative reads to peripherals is to map all storage containing code into Region 2, and all I/O devices into Region 1. Thus, accesses to Region 2 would only be for code and program data. Speculative fetches occurring in Region 2 would never access addresses in Region 1. Note that this hardware organization eliminates the need to use of the G storage attribute to protect Region 1. However, Region 1 could be set as guarded with no performance penalty, because there is no code to execute or variable data to access in Region 1.

The use of these regions could be reversed (code in Region 1 and I/O devices in Region 2), if Region 2 is set as guarded. Prefetching from the highest addresses of Region 1 could cause an attempt to speculatively access the bottom of Region 2, but guarding prevents this from occurring. The performance penalty is slight, under the assumption that code infrequently executes the instructions in the highest addresses of Region 1.

2.8.3 Summary

Software should take the following actions to prevent speculative accesses to sensitive data areas, if the sensitive data areas are not in guarded storage:

- Protect against accesses to “random” values in the LR or CTR on **blr** or **bctr** branches following **rfi**, **rfdi**, or **sc** instructions by putting appropriate instructions before or after the **rfi**, **rfdi**, or **sc** instruction. See “Fetching Past an Interrupt-Causing or Interrupt-Returning Instruction” on page 2-28.
- Protect against “running past” the end of memory into a bordering I/O device by putting an unconditional branch at the end of the memory area. See “Fetching Past an Unconditional Branch” on page 2-29.
- Recognize that a maximum of 19 words can be prefetched past an unresolved conditional branch, either down the target path or the sequential path. See “Prefetch Distance Down an Unresolved Branch Path” on page 2-28.

Of course, software should not code branches with known unsafe targets (either relative to the instruction counter, or to addresses contained in the LR or CTR), on the assumption that the targets are “protected” by code guaranteeing that the unsafe direction is not taken. The fetcher assumes that if a branch is predicted to be taken, it is safe to fetch down the target path.

2.9 Privileged Mode Operation

In the PowerPC Architecture, several terms describe two operating modes that have different instruction execution privileges. When a processor is in “privileged mode,” it can execute all instructions in the instruction set. This mode is also called the “supervisor state.” The other mode, in

which certain instructions cannot be executed, is called the “user mode,” or “problem state.” These terms are used in pairs:

Privileged	Non-privileged
Privileged Mode	User Mode
Supervisor State	Problem State

The architecture uses MSR[PR] to control the execution mode. When MSR[PR] = 1, the processor is in user mode (problem state); when MSR[PR] = 0, the processor is in privileged mode (supervisor state).

After a reset, MSR[PR] = 0.

2.9.1 MSR Bits and Exception Handling

The current value of MSR[PR] is saved, along with all other MSR bits, in the SRR1 (for non-critical interrupts) or SRR3 (for critical interrupts) upon any interrupt, and MSR[PR] is set to 0. Therefore, all exception handlers operate in privileged mode.

Attempting to execute a privileged instruction while in user mode causes a privileged violation program exception (see “Program Interrupt” on page 5-20). The PPC405 core does not execute the instruction, and the program counter is loaded with EVPR[0:15] || 0x0700, the address of an exception processing routine.

The PRR field of the Exception Syndrome Register (ESR) is set when an interrupt was caused by a privileged instruction program exception. Software is not required to clear ESR[PPR].

2.9.2 Privileged Instructions

The instructions listed in Table 2-9 are privileged and cannot be executed while in user mode (MSR[PR] = 1).

Table 2-9. Privileged Instructions

dcbi	
dccci	
dcread	
iccci	
icread	
mfocr	
mfmsr	
mfmsr	For all SPRs except CTR, LR, SPRG4–SPRG7, and XER. See “Privileged SPRs” on page 2-32
mtocr	
mtmsr	
mtmsr	For all SPRs except CTR, LR, XER. See “Privileged SPRs” on page 2-32
rfci	
rfi	

Table 2-9. Privileged Instructions (continued)

tlbia	
tlbre	
tlbsx	
tlbsync	
tlbwe	
wrtee	
wrteei	

2.9.3 Privileged SPRs

All SPRs are privileged, except for the LR, the CTR, the XER, USPRG0, and read access to SPRG4–SPRG7. Reading from the time base registers Time Base Lower (TBL) and Time Base Upper (TBU) is not privileged. These registers are read using the **mftb** instruction, rather than the **mf spr** instruction. TBL and TBU are written (with different addresses) using **mtspr**, which is privileged for these registers. Except for moves to and from non-privileged SPRs, attempts to execute **mf spr** and **mtspr** instructions while in user mode result in privileged violation program exceptions.

In a **mf spr** or **mtspr** instruction, the 10-bit SPRN field specifies the SPR number of the source or destination SPR. The SPRN field contains two five-bit subfields, SPRN_{0:4} and SPRN_{5:9}. The assembler handles the unusual register number encoding to generate the SPRF field. In the *machine code* for the **mf spr** and **mtspr** instructions, the SPRN subfields are *reversed* (ending up as SPRF_{5:9} and SPRF_{0:4}) for compatibility with the POWER Architecture.

In the PowerPC Architecture, SPR numbers having a 1 in the most-significant bit of the SPRF field are privileged.

The following example illustrates how SPR numbers appear in assembler language coding and in machine coding of the **mf spr** and **mtspr** instructions.

In assembler language coding, SRR0 is SPR 26. Note that the assembler handles the unusual register number encoding to generate the SPRF field.

```
mf spr r5,26
```

When the SPR number is considered as a binary number (0b0000011010), the most-significant bit is 0. However, the machine code for the instruction reverses the subfields, resulting in the following SPRF field: 0b1101000000. The most-significant bit is 1; SRR0 is privileged.

When an SPR number is considered as a hexadecimal number, the second digit of the three-digit hexadecimal number indicates whether an SPR is privileged. If the second digit is odd (1, 3, 5, 7, 9, B, D, F), the SPR is privileged.

For example, the SPR number of SRR0 is 26 (0x01A). The second hexadecimal digit is odd; SRR0 is privileged. In contrast, the LR is SPR 8 (0x008); the second hexadecimal digit is not odd; the LR is non-privileged.

2.9.4 Privileged DCRs

The **mt dcr** and **mf dcr** instructions themselves are privileged, in all cases. All DCRs are privileged.

2.10 Synchronization

The PPC405 core supports the synchronization operations of the PowerPC Architecture. The following book, chapter, and section numbers refer to related information in *The PowerPC Architecture: A Specification for a New Family of RISC Processors*:

- Book II, Section 1.8.1, “Storage Access Ordering” and “Enforce In-order Execution of I/O”
- Book III, Section 1.7, “Synchronization”
- Book III, Chapter 7, “Synchronization Requirements for Special Registers and Lookaside Buffers”

2.10.1 Context Synchronization

The context of a program is the environment (for example, privilege and address translation) in which the program executes. Context is controlled by the content of certain registers, such as the Machine State Register (MSR), and includes the content of all GPRs and SPRs.

An instruction or event is context synchronizing if it satisfies the following requirements:

1. All instructions that *precede* a context synchronizing operation must complete in the context that existed *before* the context synchronizing operation.
2. All instructions that *follow* a context synchronizing operation must complete in the context that exists *after* the context synchronizing operation.

Such instructions and events are called “context synchronizing operations.” In the PPC405 core, these include any interrupt, except a non-recoverable instruction machine check, and the **isync**, **rftci**, **rfti**, and **sc** instructions.

However, context specifically excludes the contents of memory. A context synchronizing operation does not guarantee that subsequent instructions observe the memory context established by previous instructions. To guarantee memory access ordering in the PPC405 core, one must use either an **eieio** instruction or a **sync** instruction. Note that for the PPC405 core, the **eieio** and **sync** instructions are implemented identically. See “Storage Synchronization” on page 2-35.

The contents of DCRs are not considered as part of the processor “context” managed by a context synchronizing operation. DCRs are not part of the processor core, and are analogous to memory-mapped registers. Their context is managed in a manner similar to that of memory contents.

Finally, implementations of the PowerPC Architecture can exempt the machine check exception from context synchronization control. If the machine check exception is exempted, an instruction that *precedes* a context synchronizing operation can cause a machine check exception *after* the context synchronizing operation occurs and additional instructions have completed.

The following scenarios use pseudocode examples to illustrate these limitations of context synchronization. Subsequent text explains how software can further guarantee “storage ordering.”

1. Consider the following instruction sequence:

```
STORE non-cachable to address XYZ
isync
XYZ instruction
```

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. Consider the following instruction sequence, which assumes that a PPC405 core is part of a standard product that uses DCRs to provide bus region control:

```
STORE non-cachable to address XYZ
isync
MTDCR to change a bus region containing XYZ
```

In this sequence, there is no guarantee that the STORE will occur before the **mtdcr** changing the bus region control DCR. The STORE could fail because of a configuration error.

Consider an interrupt that changes privileged mode. An interrupt is a context synchronizing operation, because interrupts cause the MSR to be updated. The MSR is part of the processor context; the context synchronizing operation guarantees that all instructions that precede the interrupt complete using the preinterrupt value of MSR[PR], and that all instructions that follow the interrupt complete using the postinterrupt value.

Consider, on the other hand, some code that uses **mtmsr** to change the value of MSR[PR], which changes the privileged mode. In this case, the MSR is changed, changing the context. It is possible, for example, that prefetched privileged instructions expect to execute after the **mtmsr** has changed the operating mode from privileged mode to user mode. To prevent privileged instruction program exceptions, the code must execute a context synchronization operation, such as **isync**, immediately after the **mtmsr** instruction to prevent further instruction execution until the **mtmsr** completes.

eieio or **sync** can ensure that the contents of memory and DCRs are synchronized in the instruction stream. These instructions guarantee storage ordering because all memory accesses that precede **eieio** or **sync** are completed before subsequent memory accesses. Neither **eieio** nor **sync** guarantee that instruction prefetching is delayed until the **eieio** or **sync** completes. The instructions do not cause the prefetch queues to be purged and instructions to be refetched. See “Storage Synchronization” on page 2-35 for more information.

Instruction cache state is part of context. A context synchronization operation is required to guarantee instruction cache access ordering.

3. Consider the following instruction sequence, which is required for creating self-modifying code:

```
STORE    Change data cache contents
dcbst    Flush the new data cache contents to memory
sync     Guarantee that dcbst completes before subsequent instructions begin
icbi     Context changing operation; invalidates instruction cache contents.
isync    Context synchronizing operation; causes refetch using new instruction cache context
         text and new memory context, due to the previous sync.
```

If software wishes to ensure that all storage accesses are complete before executing a **mtdcr** to change a bus region (Example 2), the software must issue a **sync** after all storage accesses and before the **mtdcr**. Likewise, if the software is to ensure that all instruction fetches after the **mtdcr** use the new bank register contents, the software must issue an **isync**, after the **mtdcr** and before the first instruction that should be fetched in the new context.

isync guarantees that all subsequent instructions are fetched and executed using the context established by all previous instructions. **isync** is a context synchronizing operation; **isync** causes all subsequently prefetched instructions to be discarded and refetched.

The following example illustrates the use of **isync** with debug exceptions:

mtdbcr0	Enable an instruction address compare (IAC) 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 was necessary to guarantee that the IAC event occurs at the execution of this instruction

2.10.2 Execution Synchronization

For completeness, consider the definition of execution synchronizing as it relates to context synchronization. Execution synchronization is architecturally a subset of context synchronization.

Execution synchronization guarantees that the following requirement is met:

All instructions that *precede* an execution synchronizing operation must complete in the context that existed *before* the execution synchronizing operation.

The following requirement need not be met:

All instructions that *follow* an execution synchronizing operation must complete in the context that exists *after* the execution synchronizing operation.

Execution synchronization ensures that preceding instructions execute in the old context; subsequent instructions might execute in either the new or old context (indeterminate). The PPC405 core provides three execution synchronizing operations: the **eieio**, **mtmsr**, and **sync** instructions.

Because **mtmsr** is execution synchronizing, it guarantees that previous instructions complete using the old MSR value. (For example, using **mtmsr** to change the endian mode.) However, to guarantee that subsequent instructions use the new MSR value, we have to insert a context synchronization operation, such as **isync**.

Note that the PowerPC Architecture requires MSR[EE] (the external interrupt bit) to be, in effect, execution synchronizing: if a **mtmsr** sets MSR[EE] = 1, and an external interrupt is pending, the exception must be taken before the instruction that follows **mtmsr** is executed. However, the **mtmsr** instruction is not a context synchronizing operation, so the PPC405 core does not, for example, discard prefetched instructions and refetch. Note that the **wrtee** and **wrteei** instructions can change the value of MSR[EE], but are not execution synchronizing.

Finally, while **sync** and **eieio** are execution synchronizing, they are also more restrictive in their requirement of memory ordering. Stating that an operation is execution synchronizing does not imply storage ordering. This is an additional specific requirement of **sync** and **eieio**.

2.10.3 Storage Synchronization

The **sync** instruction guarantees that all previous storage references complete with respect to the PPC405 core before the **sync** instruction completes (therefore, before any subsequent instructions begin to execute). The **sync** instruction is execution synchronizing.

Consider the following use of **sync**:

stw Store to peripheral
 sync Wait for store to actually complete
 mtdcr Reconfigure device

The **eieio** instruction guarantees the order of storage accesses. All storage accesses that precede **eieio** complete before any storage accesses that follow the instruction, as in the following example:

stb X Store to peripheral, address X; this resets a status bit in the device
 eieio Guarantee **stb X** completes before next instruction
 lbz Y Load from peripheral, address Y; this is the status register updated by **stb X**.
 eieio was necessary, because the read and write addresses are different, but
 affect each other

The PPC405 core implements both **sync** and **eieio** identically, in the manner described above for **sync**. In the PowerPC Architecture, **sync** can function across all processors in a multiprocessor environment; **eieio** functions only within its executing processor. The PPC405 does not provide hardware support for multiprocessor memory coherency, so **sync** does not guarantee memory ordering across multiple processors.

2.11 Instruction Set

The PPC405 instruction set contains instructions defined in the PowerPC Architecture and instructions specific to the IBM PowerPC 400 family of embedded processors.

Chapter 9, “Instruction Set,” contains detailed descriptions of each instruction.

Appendix A, “Instruction Summary,” alphabetically lists each instruction and extended mnemonic and provides a short-form description. Appendix B, “Instructions by Category,” provides short-form descriptions of instructions, grouped by the instruction categories listed in Table 2-10, “PPC405 Instruction Set Summary,” on page 2-36.

Table 2-10 summarizes the PPC405 instruction set functions by categories. Instructions within each category are described in subsequent sections.

Table 2-10. PPC405 Instruction Set Summary

Storage Reference	load, store
Arithmetic	add, subtract, negate, multiply, multiply-accumulate, multiply halfword, divide
Logical	and, andc, or, orc, xor, nand, nor, xnor, sign extension, count leading zeros
Comparison	compare, compare logical, compare immediate
Branch	branch, branch conditional, branch to LR, branch to CTR
CR Logical	crand, crandc, cror, crorc, crnand, crnor, crxor, crxnor, move CR field
Rotate	rotate and insert, rotate and mask, shift left, shift right
Shift	shift left, shift right, shift right algebraic
Cache Management	invalidate, touch, zero, flush, store, read
Interrupt Control	write to external interrupt enable bit, move to/from MSR, return from interrupt, return from critical interrupt
Processor Management	system call, synchronize, trap, move to/from DCRs, move to/from SPRs, move to/from CR

2.11.1 Instructions Specific to the IBM PowerPC Embedded Environment

To support functions required in embedded real-time applications, the IBM PowerPC 400 family of embedded processors defines instructions that are not defined in the PowerPC Architecture.

Table 2-11 lists the instructions specific to IBM PowerPC embedded processors. Programs using these instructions are not portable to PowerPC implementations that are not part of the IBM PowerPC 400 family of embedded processors.

In the table, the syntax [s] indicates that the instruction has a signed form. The syntax [u] indicates that the instruction has an unsigned form. The syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-11. Implementation-specific Instructions

dccci	macchw[s][u]	mulchw[u]	mfdcr
dcread	machhw[s][u]	mulhhw[u]	mtdcr
iccci	maclhw[s][u]	mullhw[u]	rfci
icread	nmacchw[s]		tlbre
	nmachhw[s]		tlbsx[.]
	nmaclhw[s]		tlbwe
			wrtee
			wrteei

2.11.2 Storage Reference Instructions

Table 2-12 lists the PPC405 storage reference instructions. Load/store instructions transfer data between memory and the GPRs. These instructions operate on bytes, halfwords, and words. Storage reference instructions also support loading or storing multiple registers, character strings, and byte-reversed data.

In the table, the syntax “[u]” indicates that an instruction has an “update” form that updates the RA addressing register with the calculated address, and a “non-update” form. The syntax “[x]” indicates that an instruction has an “indexed” form, which forms the address 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 (included as part of the instruction word) to the contents of RA GPR.

Table 2-12. Storage Reference Instructions

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

2.11.3 Arithmetic Instructions

Arithmetic operations are performed on integer operands stored in GPRs. 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 GPRs. The result is placed in a third, operand, which is stored in a GPR. Instructions that perform operations on one operand are defined using a two-operand format; the operation is performed on the operand in a GPR and the result is placed in another GPR. Several instructions also have immediate formats in which an operand is contained in a field in the instruction word.

Most arithmetic instructions have versions that can update CR[CR0] and XER[SO, OV], based on the result of the instruction. Some arithmetic instructions also update XER[CA] implicitly. See “Condition Register (CR)” on page 2-10 and “Fixed Point Exception Register (XER)” on page 2-7 for more information.

Table 2-13 lists the PPC405 arithmetic instructions. In the table, the syntax “[o]” indicates that an instruction has an “o” form that updates XER[SO,OV], and a “non-o” form. The syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-13. Arithmetic Instructions

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

Table 2-14 lists additional arithmetic instructions for multiply-accumulate and multiply halfword operations. In the table, the syntax “[o]” indicates that an instruction has an “o” form that updates XER[SO,OV], and a “non-o” form. The syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-14. Multiply-Accumulate and Multiply Halfword Instructions

Multiply-Accumulate	Negative-Multiply-Accumulate	Multiply Halfword
macchw[o][.] macchws[o][.] macchwsu[o][.] macchwu[o][.] machhw[o][.] machhws[o][.] machhwsu[o][.] machhwu[o][.] maclhw[o][.] maclhws[o][.] maclhwsu[o][.] maclhwu[o][.]	nmacchw[o][.] nmacchws[o][.] nmachhw[o][.] nmachhws[o][.] nmaclhw[o][.] nmaclhws[o][.]	mulchw[.] mulchwu[.] mulhhw[.] mulhhwu[.] mullhw[.] mullhwu[.]

2.11.4 Logical Instructions

Table 2-15 lists the PPC405 logical instructions. In the table, the syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-15. Logical Instructions

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

2.11.5 Compare Instructions

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

Table 2-16 lists the PPC405 core compare instructions.

Table 2-16. Compare Instructions

Arithmetic	Logical
cmp cmpi	cmpl cmpli

2.11.6 Branch Instructions

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

See “Branch Processing” on page 2-24 for more information on branch operations.

Table 2-17 lists the PPC405 branch instructions. In the table, the syntax “[I]” indicates that the instruction has a “link update” form that updates LR with the address of the instruction after the branch, and a “non-link update” form. The syntax “[a]” indicates that the instruction has 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 immediate field to the address of the branch instruction).

Table 2-17. Branch Instructions

Branch
b[l][a]
bc[l][a]
bcctr[l]
bclr[l]

2.11.6.1 CR Logical Instructions

These instructions perform logical operations on a specified pair of bits in the CR, placing the result in another specified bit. These instructions can logically combine the results of several comparisons without incurring the overhead of conditional branch instructions. Software performance can significantly improve if multiple conditions are tested at once as part of a branch decision.

Table 2-18 lists the PPC405 condition register logical instructions.

Table 2-18. CR Logical Instructions

crand	crnor
crandc	cror
creqv	crorc
crnand	crxor
	mcrf

2.11.6.2 Rotate Instructions

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

Table 2-19 lists the PPC405 rotate instructions. In the table, the syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-19. Rotate Instructions

Rotate and Insert	Rotate and Mask
rlwimi[.]	rlwinm[.] rlwnm[.]

2.11.6.3 Shift Instructions

These instructions rotate operands stored in the GPRs.

Table 2-20 lists the PPC405 shift instructions. Shift right algebraic instructions implicitly update XER[CA]. In the table, the syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-20. Shift Instructions

Shift Left	Shift Right	Shift Right Algebraic
slw[.]	srw[.]	sraw[.] srawi[.]

2.11.6.4 Cache Management Instructions

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

Table 2-21 lists the PPC405 core cache management instructions.

Table 2-21. Cache Management Instructions

DCU	ICU
dcba	icbi
dcbf	icbt
dcbi	iccci
dcbst	icread
dcbt	
dcbtst	
dcbz	
dccci	
dcread	

2.11.7 Interrupt Control Instructions

mfmsr and **mtmsr** read and write data between the MSR and a GPR to enable and disable interrupts. **wrtee** and **wrteei** enable and disable external interrupts. **rfi** and **rfci** return from interrupt handlers. Table 2-22 lists the PPC405 core interrupt control instructions.

Table 2-22. Interrupt Control Instructions

mfmsr
mtmsr
rfi
rfci
wrtee
wrteei

2.11.8 TLB Management Instructions

The TLB management instructions read and write entries of the TLB array in the MMU, search the TLB array for an entry which will translate a given address, and invalidate all TLB entries. There is also an instruction for synchronizing TLB updates with other processors, but because the PPC405 core is for use in uniprocessor environments, this instruction performs no operation.

Table 2-23 lists the TLB management instructions. In the table, the syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table 2-23. TLB Management Instructions

tlbia
tlbre
tlbsx[.]
tlbsync
tlbwe

2.11.9 Processor Management Instructions

These instructions move data between the GPRs and SPRs, the CR, and DCRs in the PPC405 core, and provide traps, system calls, and synchronization controls.

Table 2-24 lists the processor management instructions in the PPC405 core.

Table 2-24. Processor Management Instructions

eieio	mcrxr	mtrcf
isync	mfcrr	mtdcr
sync	mfcdcr	mtspr
	mfsprr	sc
		tw
		twi

2.11.10 Extended Mnemonics

In addition to mnemonics for instructions supported directly by hardware, the PowerPC Architecture defines numerous *extended mnemonics*.

An extended mnemonic translates directly into the mnemonic of a hardware instruction, typically with carefully specified operands. For example, the PowerPC Architecture does not define a “shift right word immediate” instruction, because the “rotate left word immediate then AND with mask,” (**rlwinm**) instruction can accomplish the same result:

rlwinm RA,RS,32–n,n,31

However, because the required operands are not obvious, the PowerPC Architecture defines an extended mnemonic:

srwi RA,RS,n

Extended mnemonics transfer the problem of remembering complex or frequently used operand combinations to the assembler, and can more clearly reflect a programmer’s intentions. Thus, programs can be more readable.

Refer to the following chapter and appendixes for lists of the extended mnemonics:

- Chapter 9, “Instruction Set,” lists extended mnemonics under the associated hardware instruction mnemonics.
- Appendix A, “Instruction Summary,” lists extended mnemonics alphabetically, along with the hardware instruction mnemonics.
- Table B-5 in Appendix B, “Instructions by Category,” lists all extended mnemonics.

Chapter 3. Initialization

This chapter describes reset operations, the initial state of the PPC405 core after a reset, and an example of the initialization code required to begin executing application code. Initialization of external system components or system-specific chip facilities may also be performed, in addition to the basic initialization described in this chapter.

Reset operations affect the PPC405 at power on time as well as during normal operation, if programmed to do so. To understand how these operations work it is necessary to first understand the signal pins involved as well as the terminology of core, chip and system resets. Three types of reset, each with different scope, are possible in the PPC405. A core reset affects only the processor core. Chip resets affect the processor core and all on-chip peripherals. System resets affect the processor core, all on-chip peripherals, and any off-chip devices connected to the chip reset net. Only the processor core can request a core or chip reset.

The processor core can request three types of processor resets: core, chip, and system. Each type of reset can be generated by a JTAG debug tool, by the second expiration of the watchdog timer, or by writing a non-zero value to the Reset (RST) field of Debug Control Register 0 (DCR0). In Core+ASIC and system on chip (SOC) designs, reset signals from on-chip and external peripherals can initiate system resets.

Core reset	Resets the processor core, including the data cache unit (DCU) and instruction cache unit (ICU).
Chip reset	Resets the processor core, including the DCU and ICU. This type of reset is provided in the IBM PowerPC 400 Series Embedded controllers as a means of resetting on-chip peripherals, and is provided on the PPC405 for compatibility.
System reset	Resets the entire chip. The reset signal is driven active by the PPC405 during system reset.

The effects of core and chip resets on the processor core are identical. To determine which reset type occurred, the most-recent reset (MRR) field of the Debug Status Register (DBSR) can be examined.

3.1 Processor State After Reset

After a reset, the contents of the Machine State Register (MSR) and the Special Purpose Registers (SPRs) control the initial processor state. The contents of Device Control Registers (DCRs) control the initial states of on-chip devices. Chapter 10, "Register Summary," contains descriptions of the registers.

In general, the contents of SPRs are undefined after a reset. Reset initializes the minimum number of SPR fields required for allow successful instruction fetching. "Contents of Special Purpose Registers after Reset" on page 3-3 describes these initial values. System software fully configures the processor.

"Machine State Register Contents after Reset" on page 3-2 describes the MSR contents.

The MCI field of the Exception Syndrome Register (ESR) is cleared so that it can be determined if there has been a machine check during initialization, before machine check exceptions are enabled.

Two SPRs contain status on the type of reset that has occurred. The Debug Status Register (DBSR) contains the most recent reset type. The Timer Status Register (TSR) contains the most recent watchdog reset.

3.1.1 Machine State Register Contents after Reset

After all resets, all fields of the Machine State Register (MSR) contain zeros. Table 3-1 shows how this affects core operation.

Table 3-1. MSR Contents after Reset

Register	Field	Core Reset	Chip Reset	System Reset	Comment
MSR	AP	0	0	0	APU unavailable
	APE	0	0	0	Auxiliary processor exception disabled
	WE	0	0	0	Wait state disabled
	CE	0	0	0	Critical interrupts disabled
	EE	0	0	0	External interrupts disabled
	PR	0	0	0	Supervisor mode
	FP	0	0	0	Floating point unavailable
	ME	0	0	0	Machine check exceptions disabled
	FE0	0	0	0	Floating point exception disabled
	DWE	0	0	0	Debug wait mode disabled
	DE	0	0	0	Debug interrupts disabled
	FE1	0	0	0	Floating point exceptions disabled
	DR	0	0	0	Data translation disabled
	IR	0	0	0	Instruction translation disabled

3.1.2 Contents of Special Purpose Registers after Reset

In general, the contents of Special Purpose Registers (SPRs) are undefined after a core, chip, or system reset. Some SPRs retain the contents they had before a reset occurred.

Table 3-2 shows the contents of SPRs that are defined or unchanged after core, chip, and system resets.

Table 3-2. SPR Contents After Reset

Register	Bits/Fields	Core Reset	Chip Reset	System Reset	Comment
CCR0	0:31	0x00700000	0x00700000	0x00700000	Sets ICU and DCU PLB priorities
DBCR0	EDM	0	0	0	External debug mode disabled
	RST	00	00	00	No reset action.
DBCR1	0:31	0x00000000	0x00000000	0x00000000	Data compares disabled
DBSR	MRR	01	10	11	Most recent reset
DCCR	S0:S31	0x00000000	0x00000000	0x00000000	Data cache disabled
DCWR	W0:W31	0x00000000	0x00000000	0x00000000	Data cache write-through disabled
ESR	0:31	0x00000000	0x00000000	0x00000000	No exception syndromes
ICCR	S0:S31	0x00000000	0x00000000	0x00000000	Instruction cache disabled
PVR	0:31				Processor version
SGR	G0:G31	0xFFFFFFFF	0xFFFFFFFF	0xFFFFFFFF	Storage is guarded
SLER	S0:S31	0x00000000	0x00000000	0x00000000	Storage is big endian
SU0R	K0:K31	0x00000000	0x00000000	0x00000000	Storage is uncompressed
TCR	WRC	00	00	00	Watchdog timer reset disabled
TSR	WRS	Copy of TCR[WRC]	Copy of TCR[WRC]	Copy of TCR[WRC]	Watchdog reset status
	PIS	Undefined	Undefined	Undefined	After POR
	FIS	Unchanged	Unchanged	Unchanged	If reset not caused by watchdog timer

3.2 PPC405 Initial Processor Sequencing

After any reset, the processor core fetches the word at address 0xFFFFFFFFC and attempts to execute it. The instruction at 0xFFFFFFFFC is typically a branch to initialization code. Unless the instruction at 0xFFFFFFFFC is an unconditional branch, fetching can wrap to address 0x00000000 and attempt to execute the instruction at this location.

Because the processor is initially in big endian mode, initialization code must be in big endian format until the endian storage attribute for the addressed region is changed, or until code branches to a region defined as little endian storage.

Before a reset operation begins, the system must provide non-volatile memory, or memory initialized by some mechanism external to the processor. This memory must be located at address 0xFFFFF000.

3.3 Initialization Requirements

When any reset is performed, the processor is initialized to a minimum configuration to start executing initialization code. Initialization code is necessary to complete the processor and system configuration.

The initialization code example in this section performs the configuration tasks required to prepare the PPC405 core to boot an operating system or run an application program.

Some portions of the initialization code work with system components that are beyond the scope of this manual.

Initialization code should perform the following tasks to configure the processor resources.

To improve instruction fetching performance: initialize the SGR appropriately for guarded or unguarded storage. Since all storage is initially guarded and speculative fetching is inhibited to guarded storage, reprogramming the SGR will improve performance for unguarded regions.

1. Before executing instructions as cachable:
 - Invalidate the instruction cache.
 - Initialize the ICCR to configure instruction cachability.
2. Before using storage access instructions:
 - Invalidate the data cache.
 - Initialize CRRO to determine if a store miss results in a line fill (SWOA).
 - Initialize the DCWR to select copy-back or write-through caching.
 - Initialize the DCCR to configure data cachability.
3. Before allowing interrupts (synchronous or asynchronous):
 - Initialize the EVPR to point to vector table.
 - Provide vector table with branches to interrupt handlers.
4. Before enabling asynchronous interrupts:
 - Initialize timer facilities.
 - Initialize MSR to enable appropriate interrupts.
5. Initialize other processor features, such as the MMU, APU (if implemented), debug, and trace.
6. Initialize non-processor resources.
 - Initialize system memory as required by the operating system or application code.
 - Initialize off-chip system facilities.
7. Start the execution of operating system or application code.

3.4 Initialization Code Example

The following initialization code illustrates the steps that should be taken to initialize the processor before an operating system or user programs begin execution. The example is presented in pseudo-code; function calls are named similarly to PPC405 mnemonics where appropriate. Specific implementations may require different ordering of these sections to ensure proper operation.

```

/* _____ */
/*  PPC405 Initialization Pseudo Code          */
/* _____ */
@0xFFFFFFFF:          /* initial instruction fetch from 0xFFFFFFFF */
    ba(init_code);    /* branch to initialization code */

@init_code:

/* _____ */
/* Configure guarded attribute for performance. */
/* _____ */

    mtspr(SGR, guarded_attribute);

/* _____ */
/* Configure endianness and compression.      */
/* _____ */

    mtspr(SLER, endianness);
    mtspr(SU0R, compression_attribute);

/* _____ */
/* Invalidate the instruction cache and enable cachability */
/* _____ */

iccci;                /* invalidate i-cache */
    mtspr(ICCR, i_cache_cachability);    /* enable I-cache */
isync;

/* _____ */
/* Invalidate the data cache and enable cachability */
/* _____ */

address = 0;          /* start at first line */
for (line = 0; line <m_lines; line++) /* D-cache has m_lines congruence classes */
{
    dccci(address);    /* invalidate congruence class */
    address += 32;     /* point to the next congruence class */
}
    mtspr(CCR0, store-miss_line-fill);
    mtspr(DCWR, copy-back_write-thru);
    mtspr(DCCR, d_cache_cachability);    /* enable D-cache */
isync;

/* _____ */
/* Prepare system for synchronous interrupts. */
/* _____ */

```

```

mtspr(EVPR, prefix_addr);          /* initialize exception vector prefix */

/* Initialize vector table and interrupt handlers if not already done */

/* Initialize and configure timer facilities */

mtspr(PIT, 0);                      /* clear PIT so no PIT indication after TSR cleared*/
mtspr(TSR, 0xFFFFFFFF);           /* clear TSR */
mtspr(TCR, timer_enable);         /* enable desired timers */
mtspr(TBL, 0);                     /* reset time base low first to avoid ripple */
mtspr(TBU, time_base_u);          /* set time base, hi first to catch possible ripple */
mtspr(TBL, time_base_l);          /* set time base, low */
mtspr(PIT, pit_count);            /* set desired PIT count */

/* Initialize the MSR */

/*----- */
/* Exceptions must be enabled immediately after timer facilities to avoid missing a */
/* timer exception. */
/*----- */
/* The MSR also controls privileged/user mode, translation, and the wait state. */
/* These must be initialized by the operating system or application code. */
/* If enabling translation, code must initialize the TLB. */
/*----- */

mtmsr(machine_state);

/*----- */
/* Initialization of other processor facilities should be performed at this time. */
/*----- */

/*----- */
/* Initialization of non-processor facilities should be performed at this time. */
/*----- */

/*----- */
/* Branch to operating system or application code can occur at this time. */
/*----- */

```

Chapter 4. Cache Operations

The PPC405 core incorporates two internal cache units, an instruction cache unit (ICU) and a data cache unit (DCU). Instructions and data can be accessed in the caches much faster than in main memory, if instruction and data cache arrays are implemented. The PPC405B3 core has a 16KB instruction cache array and an 8KB data cache array.

The ICU controls instruction accesses to main memory and, if an instruction cache array is implemented, stores frequently used instructions to reduce the overhead of instruction transfers between the instruction pipeline and external memory. Using the instruction cache minimizes access latency for frequently executed instructions.

The DCU controls data accesses to main memory and, if a data cache array is implemented, stores frequently used data to reduce the overhead of data transfers between the GPRs and external memory. Using the data cache minimizes access latency for frequently used data.

The ICU features:

- Programmable address pipelining and prefetching for cache misses and non-cachable lines
- Support for non-cachable hits from lines contained in the line fill buffer
- Programmable non-cachable requests to memory as 4 or 8 words (or half line or line)
- Bypass path for critical words
- Non-blocking cache for hits during fills
- Flash invalidate (one instruction invalidates entire cache)
- Programmable allocation for fetch fills, enabling program control of cache contents using the **icbt** instruction
- Virtually indexed, physically tagged cache arrays
- A rich set of cache control instructions

The DCU features:

- Address pipelining for line fills
- Support for load hits from non-cachable and non-allocated lines contained in the line fill buffer
- Bypass path for critical words
- Non-blocking cache for hits during fills
- Write-back and write-through write strategies controlled by storage attributes
- Programmable non-cachable load requests to memory as lines or words.
- Handling of up to two pending line flushes.
- Holding of up to three stores before stalling the core pipeline
- Physically indexed, physically tagged cache arrays
- A rich set of cache control instructions

The PPC405 core can include an instruction cache array and a data cache array. The size of the cache arrays can vary by core implementation, as shown in Table 4-1.

Table 4-1. Available Cache Array Sizes

ICU Cache Array Size	DCU Cache Array Size
0KB	0KB
4KB	4KB
8KB	8KB
16KB	16KB
32KB	32KB

Programming Note: If the ICU cache array or the DCU cache array is not present (0KB), the I (cachability) storage attribute must be turned off for instruction-side or data-side memory, respectively.

“ICU and DCU Organization and Sizes” describes the organization and sizes of the ICU and the DCU. “ICU Overview” on page 4-3 and “DCU Overview” on page 4-6 provide overviews of the ICU and DCU.

4.1 ICU and DCU Organization and Sizes

The ICU and DCU contain control logic and, in some implementations, cache arrays. The control logic, which handles data transfers between the cache units, main memory, and the RISC core, differs significantly between the ICU and DCU. The ICU and DCU cache arrays, which (when implemented) store instructions and data from main memory, respectively, are almost identical. (The DCU array adds a “dirty” bit to mark modified lines.)

The ICU and DCU cache arrays are two-way set-associative. In both cache units, a cache line can be in one of two locations in the cache array. The two locations are members of a set of locations. Each set is divided into two ways, way A and way B; a cache line can be located in either way. Each way is organized as n lines of eight words each, where n is the cache size, in kilobytes, multiplied by 16. For example, a 4KB cache array contains 64 lines.

Cache lines are addressed using a tag field and an index. The tag fields are also two-way set-associative. As shown in Table 4-2, the tag fields in ways A and B store address bits $A_{0:21}$ for each

cache line. The remaining address bits ($A_{22:27}$) serve as an index to the cache array. The two cache lines that correspond with the same line index are called a congruence class.

Table 4-2. ICU and DCU Cache Array Organization

Tags (Two-way Set)		Cache Lines (Two-way Set)	
Way A	Way B	Way A	Way B
$A_{0:m-1}$ Line 0	$A_{0:m-1}$ Line 0	Line 0	Line 0
$A_{0:m-1}$ Line 1	$A_{0:m-1}$ Line 1	Line 1	Line 1
•	•	•	•
•	•	•	•
•	•	•	•
$A_{0:m-1}$ Line $n-2$	$A_{0:m-1}$ Line $n-2$	Line $n-2$	Line $n-2$
$A_{0:m-1}$ Line $n-1$	$A_{0:m-1}$ Line $n-1$	Line $n-1$	Line $n-1$

Table 4-3 shows the values of m and n for various cache array sizes.

Table 4-3. Cache Sizes, Tag Fields, and Lines

Array Size	Instruction Cache Array		Data Cache Array	
	m (Tag Field Bits)	n (Lines)	m (Tag Field Bits)	n (Lines)
0KB	—	—	—	—
4KB	22 (0:21)	64	20 (0:19)	64
8KB	22 (0:21)	128	20 (0:19)	128
16KB	22 (0:21)	256	20 (0:19)	256
32KB	22 (0:21)	512	20 (0:19)	512

When the ICU or DCU requests a cache line from main memory (an operation called a cache line fill), a least-recently-used (LRU) policy determines which cache line way will receive the requested line. The index, determined by the instruction or data address, selects a congruence class. Within a congruence class, the most recently accessed line (in either way A or way B) is retained and the LRU bit in the associated tag array marks the other line as LRU. The LRU line then receives the requested instruction or data words. After the cache line fill, the LRU bit is set to identify as LRU the line opposite the line just filled.

4.2 ICU Overview

The ICU manages instruction transfers between external cachable memory and the instruction queue in the execution unit.

Figure 4-1 shows the relationships between the ICU and the instruction pipeline.




Figure 4-1. Instruction Flow

4.2.1 ICU Operations

Instructions from cachable memory regions are copied into the instruction cache array, if an array is present. The fetcher can access instructions much more quickly from a cache array than from memory. Cache lines can be loaded either target-word-first or sequentially, or in any order. Target-word-first fills start at the requested word, continue to the end of the line, and then wrap to fill the remaining words at the beginning of the line. Sequential fills start at the first word of the cache line and proceed sequentially to the last word of the line.

The bypass path handles instructions in cache-inhibited memory and improves performance during line fill operations. If a request from the fetcher obtains an entire line from memory, the queue does not have to wait for the entire line to reach the cache. The target word (the word requested by the fetcher) is sent on the bypass path to the queue while the line fill proceeds, even if the selected line fill order is not target-word-first.

Cache line fills always run to completion, even if the instruction stream branches away from the rest of the line. As requested instructions are received, they go to the fetcher from the fill register before the line fills in the cache. The filled line is always placed in the ICU; if an external memory subsystem error occurs during the fill, the line is not written to the cache. During a clock cycle, the ICU can send two instructions to the fetcher.

4.2.2 Instruction Cachability Control

When instruction address translation is enabled ($MSR[IR] = 1$), instruction cachability is controlled by the I storage attribute in the translation lookaside buffer (TLB) entry for the memory page. If $TLB_entry[I] = 1$, caching is inhibited; otherwise caching is enabled. Cachability is controlled separately for each page, which can range in size from 1KB to 16MB. "Translation Lookaside Buffer (TLB)" on page 7-2 describes the TLB.

When instruction address translation is disabled ($MSR[IR] = 0$), instruction cachability is controlled by the Instruction Cache Cachability Register (ICCR). Each field in the ICCR ($ICCR[S0:S31]$) controls the cachability of a 128MB region (see "Real-Mode Storage Attribute Control" on page 7-17). If $ICCR[S_n] = 1$, caching is enabled for the specified region; otherwise, caching is inhibited.

The performance of the PPC405 core is significantly lower while fetching instructions from cache-inhibited regions.

Following system reset, address translation is disabled and all ICCR bits are reset to 0 so that no memory regions are cachable. Before regions can be designated as cachable, the ICU cache array must be invalidated, if an array is present. The **iccci** instruction must execute before the cache is enabled. Address translation can then be enabled, if required, and the TLB or the ICCR can then be configured for the required cachability.

4.2.3 Instruction Cache Synonyms

The following information applies only if instruction address translation is enabled ($MSR[IR] = 1$) and 1KB or 4KB page sizes are used. See Chapter 7, "Memory Management," for information about address translation and page sizes.

An instruction cache synonym occurs when the instruction cache array contains multiple cache lines from the same real address. Such synonyms result from combinations of:

- Cache array size
- Cache associativity
- Page size
- The use of effective addresses (EAs) to index the cache array

For example, the instruction cache array has a "way size" of 8KB (16KB array/2 ways). Thus, 11 bits ($EA_{19:29}$) are needed to select a word (instruction) in each way. For the minimum page size of 1KB, the low order 8 bits ($EA_{22:29}$) address a word in a page. The high order address bits ($EA_{0:21}$) are translated to form a real address (RA), which the ICU uses to perform the cache tag match. Cache synonyms could occur because the index bits ($EA_{19:29}$) overlap the translated RA bits. For 1KB pages, overlap in $EA_{19:21}$ and $RA_{19:21}$ could result in as many as 8 synonyms. In other words, data from the same RA could occur as many as 8 locations in the cache array. Similarly, for 4KB pages, $EA_{0:19}$ are translated. Differences in EA_{19} and RA_{19} could result in as many as 2 synonyms. For the next largest page size (16KB), only $EA_{0:17}$ are translated. Because there is no overlap with index bits $EA_{19:21}$, synonyms do not occur.

In practice, cache synonyms occur when a real instruction page having multiple virtual mappings exists in multiple cache lines. For 1KB pages, all EAs differing in EA_{19:21} must be cast out of cache, using an **icbi** instruction for each such EA (up to 8 per cache line in the page). For 4KB pages, all EAs differing in EA₁₉ must be cast out in the same manner (up to 2 per cache line in the page). For larger pages, cache synonyms do not occur, and casting out any of the multiple EAs removes the physical information from the cache.

Programming Note: To prevent the occurrence of cache synonyms, use only page sizes greater than the cache way size (8KB), if possible. For the PPC405, the minimum such page size is 16KB.

4.2.4 ICU Coherency

The ICU does not “snoop” external memory or the DCU. Programmers must follow special procedures for ICU synchronization when self-modifying code is used or if a peripheral device updates memory containing instructions.

The following code example illustrates the necessary steps for self-modifying code. This example assumes that *addr1* is both data and instruction cachable.

```

stw      regN, addr1  # the data in regN is to become an instruction at addr1
dcbst   addr1        # forces data from the data cache to memory
sync    # wait until the data actually reaches the memory
icbi    addr1        # the previous value at addr1 might already be in
                    # the instruction cache; invalidate it in the cache
isync   # the previous value at addr1 may already have been
                    # pre-fetched into the queue; invalidate the queue
                    # so that the instruction must be re-fetched

```

4.3 DCU Overview

The DCU manages data transfers between external cachable memory and the general-purpose registers in the execution unit.

A bypass path handles data operations in cache-inhibited memory and improves performance during line fill operations.

4.3.1 DCU Operations

Data from cachable memory regions are copied from external memory into lines in the data cache array so that subsequent cache operations result in cache hits. Loads and stores that hit in the DCU are completed in one cycle. For loads, GPRs receive the requested byte, halfword, or word of data from the data cache array. The DCU supports byte-writeability to improve the performance of byte and halfword store operations.

Cache operations require a line fill when they require data from cachable memory regions that are not currently in the DCU. A line fill is the movement of a cache line (eight words) from external memory to the data cache array. Eight words are copied from external memory into the fill buffer, either target-word-first or sequentially, or in any other order. Loading order is controlled by the PLB slave. Target-word-first fills start at the requested word, continue to the end of the line, and then wrap to fill the remaining words at the beginning of the line. Sequential fills start at the first word of the cache line

and proceed sequentially to the last word of the line. In both types of fills, the fill buffer, when full, is transferred to the data cache array. The cache line is marked valid when it is filled.

Loads that result in a line fill, and loads from non-cachable memory, are sent to a GPR. The requested byte, halfword, or word is sent from the DCU to the GPR from the fill buffer, using a cache bypass mechanism. Additional loads for data in the fill buffer can be bypassed to the GPR until the data is moved into the data array.

Stores that result in a line fill have their data held in the fill buffer until the line fill completes. Additional stores to the line being filled will also have their data placed in the fill buffer before being transferred into the data cache array.

To complete a line fill, the DCU must access the tag and data arrays. The tag array is read to determine the tag addresses, the LRU line, and whether the LRU line is dirty. A dirty cache line is one that was accessed by a store instruction after the line was established, and can be inconsistent with external memory. If the line being replaced is dirty, the address and the cache line must be saved so that external memory can be updated. During the cache line fill, the LRU bit is set to identify the line opposite the line just filled as LRU.

When a line fill completes and replaces a dirty line, a line flush begins. A flush copies updated data in the data cache array to main storage. Cache flushes are always sequential, starting at the first word of the cache line and proceeding sequentially to the end of the line.

Cache lines are always completely flushed or filled, even if the program does not request the rest of the bytes in the line, or if a bus error occurs after a bus interface unit accepts the request for the line fill. If a bus error occurs during a line fill, the line is filled and the data is marked valid. However, the line can contain invalid data, and a machine check exception occurs.

4.3.2 DCU Write Strategies

DCU operations can use write-back or write-through strategies to maintain coherency with external cachable memory.

The write-back strategy updates only the data cache, not external memory, during store operations. Only modified data lines are flushed to external memory, and then only when necessary to free up locations for incoming lines, or when lines are explicitly flushed using **dcbf** or **dcbst** instructions. The write-back strategy minimizes the amount of external bus activity and avoids unnecessary contention for the external bus between the ICU and the DCU.

The write-back strategy is contrasted with the write-through strategy, in which stores are written simultaneously to the cache and to external memory. A write-through strategy can simplify maintaining coherency between cache and memory.

When data address translation is enabled ($MSR[DR] = 1$), the *W* storage attribute in the TLB entry for the memory page controls the write strategy for the page. If $TLB_entry[W] = 0$, write-back is selected; otherwise, write-through is selected. The write strategy is controlled separately for each page. “Translation Lookaside Buffer (TLB)” on page 7-2 describes the TLB.

When data address translation is disabled ($MSR[DR] = 0$), the Data Cache Write-through Register (DCWR) sets the storage attribute. Each bit in the DCWR ($DCWR[W0:W31]$) controls the write strategy of a 128MB storage region (see “Real-Mode Storage Attribute Control” on page 7-17). If $DCWR[Wn] = 0$, write-back is enabled for the specified region; otherwise, write-through is enabled.

Programming Note: The PowerPC Architecture does not support memory models in which

write-through is enabled and caching is inhibited.

4.3.3 DCU Load and Store Strategies

The DCU can control whether a load receives one word or one line of data from main memory.

For cachable memory, the load without allocate (LWOA) field of the CCR0 controls the type of load resulting from a load miss. If CCR0[LWOA] = 0, a load miss causes a line fill. If CCR0[LWOA] = 1, load misses do not result in a line fill, but in a word load from external memory. For infrequent reads of non-contiguous memory, setting CCR0[LWOA] = 1 may provide a small performance improvement.

For non-cachable memory and for loads misses when CCR0[LWOA] = 1, the load word as line (LWL) field in the CCR0 affects whether load misses are satisfied with a word, or with eight words (the equivalent of a cache line) of data. If CCR0[LWL] = 0, only the target word is bypassed to the core. If CCR0[LWL] = 1, the DCU saves eight words (one of which is the target word) in the fill buffer and bypasses the target data to the core to satisfy the load word request. The fill buffer is not written to the data cache array.

Setting CCR0[LWL] = 1 provides the fastest accesses to sequential non-cachable memory. Subsequent loads from the same line are bypassed to the core from the fill buffer and do not result in additional external memory accesses. The load data remains valid in the fill buffer until one of the following occurs: the beginning of a subsequent load that requires the fill buffer, a store to the target address, a **dcbi** or **dccci** instruction issued to the target address, or the execution of a **sync** instruction. Non-cachable loads to guarded storage never cause a line transfer on the PLB even if CCR0[LWL] = 1. Subsequent loads to the same non-cachable storage are always requested again from the PLB.

For cachable memory, the store without allocate (SWOA) field of the CCR0 controls the type of store resulting from a store miss. If CCR0[SWOA] = 0, a store miss causes a line fill. If CCR0[SWOA] = 1, store misses do not result in a line fill, but in a single word store to external memory.

4.3.4 Data Cachability Control

When data address translation is disabled (MSR[DR] = 0), data cachability is controlled by the Data Cache Cachability Register (DCCR). Each bit in the DCCR (DCCR[S0:S31]) controls the cachability of a 128MB region (see “Real-Mode Storage Attribute Control” on page 7-17). If DCCR[S n] = 1, caching is enabled for the specified region; otherwise, caching is inhibited.

When data address translation is enabled (MSR[DR] = 1), data cachability is controlled by the I bit in the TLB entry for the memory page. If TLB_entry[I] = 1, caching is inhibited; otherwise caching is enabled. Cachability is controlled separately for each page, which can range in size from 1KB to 16MB. “Translation Lookaside Buffer (TLB)” on page 7-2 describes the TLB.

Programming Note: The PowerPC Architecture does not support memory models in which write-through is enabled and caching is inhibited.

The performance of the PPC405 core is significantly lower while accessing memory in cache-inhibited regions.

Following system reset, address translation is disabled and all DCCR bits are reset to 0 so that no memory regions are cachable. If an array is present, the **dccci** instruction must execute n times before regions can be designated as cachable. This invalidates all congruence classes before

enabling the cache. Address translation can then be enabled, if required, and the TLB or the DCCR can then be configured for the desired cachability.

Programming Note: If a data block corresponding to the effective address (EA) exists in the cache, but the EA is non-cachable, loads and stores (including **dcbz**) to that address are considered programming errors (the cache block should previously have been flushed). The only instructions that can legitimately access such an EA in the data cache are the cache management instructions **dcbf**, **dcbi**, **dcbst**, **dcbt**, **dcbtst**, **dccci**, and **dcread**.

4.3.5 DCU Coherency

The DCU does not provide snooping. Application programs must carefully use cache-inhibited regions and cache control instructions to ensure proper operation of the cache in systems where external devices can update memory.

4.4 Cache Instructions

For detailed descriptions of the instructions described in the following sections, see Chapter 9, “Instruction Set.”

In the instruction descriptions, the term “block” is synonymous with cache line. A block is the unit of storage operated on by all cache block instructions.

4.4.1 ICU Instructions

The following instructions control instruction cache operations:

icbi	Instruction Cache Block Invalidate Invalidates a cache block.
icbt	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. This is a privileged instruction.
iccci	Instruction Cache Congruence Class Invalidate Invalidates the instruction cache array. This is a privileged instruction.
icread	Instruction Cache Read Reads either an instruction cache tag entry or an instruction word from an instruction cache line, typically for debugging. Fields in CCR0 control instruction behavior (see “Cache Control and Debugging Features” on page 4-11). This is a privileged instruction.

4.4.2 DCU Instructions

Data cache flushes and fills are triggered by load, store and cache control instructions. Cache control instructions are provided to fill, flush, or invalidate cache blocks.

The following instructions control data cache operations.

- dcba** Data Cache Block Allocate
- Speculatively establishes a line in the cache and marks the line as modified.
- If the line is not currently in the cache, the line is established and marked as modified without actually filling the line from external memory.
- If **dcba** references a non-cachable address, **dcba** is treated as a no-op.
- If **dcba** references a cachable address, write-through required (which would otherwise cause an alignment exception), **dcba** is treated as a no-op.
- dcbf** Data Cache Block Flush
- Flushes a line, if found in the cache and marked as modified, to external memory; the line is then marked invalid.
- If the line is found in the cache and is not marked modified, the line is marked invalid but is not flushed.
- This operation is performed regardless of whether the address is marked cachable.
- dcbi** Data Cache Block Invalidate
- Invalidates a block, if found in the cache, regardless of whether the address is marked cachable. Any modified data is not flushed to memory.
- This is a privileged instruction.
- dcbst** Data Cache Block Store
- Stores a block, if found in the cache and marked as modified, into external memory; the block is not invalidated but is no longer marked as modified.
- If the block is marked as not modified in the cache, no operation is performed.
- This operation is performed regardless of whether the address is marked cachable.
- dcbt** Data Cache Block Touch
- Fills a block with data, if the address is cachable and the data is not already in the cache. If the address is non-cachable, this instruction is a no-op.
- dcbstst** Data Cache Block Touch for Store
- Implemented identically to the **dcbt** instruction for compatibility with compilers and other tools.

- dcbz** Data Cache Block Set to Zero
 Fills a line in the cache with zeros and marks the line as modified.
 If the line is not currently in the cache (and the address is marked as cachable and non-write-through), the line is established, filled with zeros, and marked as modified without actually filling the line from external memory. If the line is marked as either non-cachable or write-through, an alignment exception results.
- dccci** Data Cache Congruence Class Invalidate
 Invalidates a congruence class (both cache ways).
 This is a privileged instruction.
- dcread** Data Cache Read
 Reads either a data cache tag entry or a data word from a data cache line, typically for debugging. Bits in CCR0 control instruction behavior (see “Cache Control and Debugging Features” on page 4-11).
 This is a privileged instruction.

4.5 Cache Control and Debugging Features

Registers and instructions are provided to control cache operation and help debug cache problems. For ICU debug, the **icread** instruction and the Instruction Cache Debug Data Register (ICDBDR) are provided. See “ICU Debugging” on page 4-14 for more information. For DCU debug, the **dcread** instruction is provided. See “DCU Debugging” on page 4-15 for more information.

CCR0 controls the behavior of the **icread** and the **dcread** instructions.




Figure 4-2. Core Configuration Register 0 (CCR0)

0:5		Reserved
6	LWL	Load Word as Line 0 The DCU performs load misses or non-cachable loads as words, halfwords, or bytes, as requested 1 For load misses or non-cachable loads, the DCU moves eight words (including the target word) into the line fill buffer
7	LWOA	Load Without Allocate 0 Load misses result in line fills 1 Load misses do not result in a line fill, but in non-cachable loads

8	SWOA	Store Without Allocate 0 Store misses result in line fills 1 Store misses do not result in line fills, but in non-cachable stores	
9	DPP1	DCU PLB Priority Bit 1 0 DCU PLB priority 0 on bit 1 1 DCU PLB priority 1 on bit 1	Note: DCU logic dynamically controls DCU priority bit 0.
10:11	IPP	ICU PLB Priority Bits 0:1 00 Lowest ICU PLB priority 01 Next to lowest ICU PLB priority 10 Next to highest ICU PLB priority 11 Highest ICU PLB priority	
12:13		Reserved	
14	U0XE	Enable U0 Exception 0 Disables the U0 exception 1 Enables the U0 exception	
15	LDBE	Load Debug Enable 0 Load data is invisible on data-side (on-chip memory (OCM)) 1 Load data is visible on data-side OCM	
16:19		Reserved	
20	PFC	ICU Prefetching for Cachable Regions 0 Disables prefetching for cachable regions 1 Enables prefetching for cachable regions	
21	PFNC	ICU Prefetching for Non-Cachable Regions 0 Disables prefetching for non-cachable regions 1 Enables prefetching for non-cachable regions	
22	NCRS	Non-cachable ICU request size 0 Requests are for four-word lines 1 Requests are for eight-word lines	
23	FWOA	Fetch Without Allocate 0 An ICU miss results in a line fill. 1 An ICU miss does not cause a line fill, but results in a non-cachable fetch.	
24:26		Reserved	
27	CIS	Cache Information Select 0 Information is cache data. 1 Information is cache tag.	
28:30		Reserved	
31	CWS	Cache Way Select 0 Cache way is A. 1 Cache way is B.	

4.5.1 CCR0 Programming Guidelines

Several fields in CCR0 affect ICU and DCU operation. Altering these fields while the cache units are involved in PLB transfers can cause errant operation, including a processor hang.

To guarantee correct ICU and DCU operation, specific code sequences must be followed when altering CCR0 fields.

CCR0[IPP, FWOA] affect ICU operation. When these fields are altered, execution of the following code sequence (Sequence 1) is required.

```
! SEQUENCE 1 Altering CCR0[IPP, FWOA]
! Turn off interrupts
mfmsr    RM
addis    RZ,r0,0x0002 ! CE bit
ori      RZ,RZ,0x8000 ! EE bit
andc     RZ,RM,RZ    ! Turn off MSR[CE,EE]
mtmsr    RZ
! sync
sync
! Touch code sequence into i-cache
addis    RX,r0,seq1@h
ori      RX,RX,seq1@l
icbt     r0,RX
! Call function to alter CCR0 bits
    b seq1
back:
! Restore MSR to original value
    mtmsr    RM
        .
        .
        .
! The following function must be in cacheable memory
    .align 5    ! Align CCR0 altering code on a cache line boundary.
seq1:
icbt     r0,RX          ! Repeat ICBT and execute an ISYNC to guarantee CCR0
isync                    ! altering code has been completely fetched across the PLB.
mfspr    RN,CCR0      ! Read CCR0.
andi/ori RN,RN,0XXXXX ! Execute and/or function to change any CCR0 bits.
                        ! Can use two instructions before having to touch
                        ! in two cache lines.
mtspr    CCR0, RN ! Update CCR0.
isync                    ! Refetch instructions under new processor context.
b    back    ! Branch back to initialization code.
```

CCR0[DPP1, U0XE] affect DCU operation. When these fields are altered, execution of the following code sequence (Sequence 2) is required. Note that Sequence 1 includes Sequence 2, so Sequence 1 can be used to alter any CCR0 fields.

In the following sample code, registers RN, RM, RX, and RZ are any available GPRs.

```

! SEQUENCE 2 Alter CCR0[DPP1, U0XE)
! Turn off interrupts
    mfmsr    RM
    addis    RZ,r0,0x0002 ! CE bit
    ori      RZ,RZ,0x8000 ! EE bit
    andc     RZ,RM,RZ     ! Turn off MSR[CE,EE]
    mtmsr    RZ
! sync
    sync
! Alter CCR0 bits
    mfspr    RN,CCR0     ! Read CCR0.
    andi/ori RN,RN,0xFFFF ! Execute and/or function to change any CCR0 bits.
    mtspr    CCR0, RN    ! Update CCR0.
    isync    ! Refetch instructions under new processor context.
! Restore MSR to original value
    mtmsr    RM

```

CCR0[CIS, CWS] do not require special programming.

4.5.2 ICU Debugging

The **icread** instruction enables the reading of the instruction cache entries for the congruence class specified by EA_{18:26}, unless no cache array is present. The cache information is read into the ICDBDR; from there it can subsequently be moved, using a **mfspr** instruction, into a GPR.

0	31
---	----

Figure 4-3. Instruction Cache Debug Data Register (ICDBDR)

0:31	Instruction cache information	See icread , page -68.
------	-------------------------------	-------------------------------

ICU tag information is placed into the ICDBDR as shown:

0:21	TAG	Cache Tag
22:26		Reserved
27	V	Cache Line Valid 0 Not valid 1 Valid
28:30		Reserved
31	LRU	Least Recently Used (LRU) 0 A-way LRU 1 B-way LRU

If CCR0[CIS] = 0, the data is a word of ICU data from the addressed line, specified by EA_{27:29}. If CCR0[CWS] = 0, the data is from the A-way; otherwise; the data from the B-way.

If CCR0[CIS] = 1, the cache information is the cache tag. If CCR0[CWS] = 0, the tag is from the A-way; otherwise, the tag is from the B-way.

Programming Note: The instruction pipeline does not wait for data from an **icread** instruction to arrive before attempting to use the contents the ICDBDR. The following code sequence ensures proper results:

```
icread r5,r6# read cache information
isync      # ensure completion of icread
mficbdr r7# move information to GPR
```

4.5.3 DCU Debugging

The **dcread** instruction provides a debugging tool for reading the data cache entries for the congruence class specified by EA_{18:26}, unless no cache array is present. The cache information is read into a GPR.

If CCR0[CIS] = 0, the data is a word of DCU data from the addressed line, specified by EA_{27:29}. If EA_{30:31} are not 00, an alignment exception occurs. If CCR0[CWS] = 0, the data is from the A-way; otherwise; the data is from the B-way.

If CCR0[CIS] = 1, the cache information is the cache tag. If CCR0[CWS] = 0, the tag is from the A-way; otherwise the tag is from the B-way.

DCU tag information is placed into the GPR as shown:

0:19	TAG	Cache Tag
20:25		Reserved
26	D	Cache Line Dirty 0 Not dirty 1 Dirty
27	V	Cache Line Valid 0 Not valid 1 Valid
28:30		Reserved
31	LRU	Least Recently Used (LRU) 0 A-way LRU 1 B-way LRU

Note: A “dirty” cache line is one which has been accessed by a store instruction after it was established, and can be inconsistent with external memory.

4.6 DCU Performance

DCU performance depends upon the application and the design of the attached external bus controller, but, in general, cache hits complete in one cycle without stalling the CPU pipeline. Under certain conditions and limitations of the DCU, the pipeline stalls (stops executing instructions) until the DCU completes current operations.

Several factors affect DCU performance, including:

- Pipeline stalls
- DCU priority
- Simultaneous cache operations
- Sequential cache operations

4.6.1 Pipeline Stalls

The CPU issues commands for cache operations to the DCU. If the DCU can immediately perform the requested cache operation, no pipeline stall occurs. In some cases, however, the DCU cannot immediately perform the requested cache operation, and the pipeline stalls until the DCU can perform the pending cache operation.

In general, the DCU, when hitting in the cache array, can execute a load/store every cycle. If a cache miss occurs, the DCU must retrieve the line from main memory. For cache misses, the DCU stores the cache line in a line fill buffer until the entire cache line is received. The DCU can accept new DCU commands while the fill progresses. If the instruction causing the line fill is a load, the target word is bypassed to the GPR during the cycle after it becomes available in the fill buffer. When the fill buffer is full, it must be moved into the tag and data arrays. During this time, the DCU cannot begin a new cache operation and stalls the pipeline if new DCU commands are presented. Storing a line in the line fill buffer takes 3 cycles, unless the line being replaced has been modified. In that case, the operation takes 4 cycles.

The DCU can accept up to two load commands. If the data for the first load command is not immediately available, the DCU can still accept the second load command. If the load data is not required by subsequent instructions, those instructions will continue to execute. If data is required from either load command, the CPU pipeline will stall until the load data has been delivered. The pipeline will also stall until the second load has read the data array if a subsequent data cache command is issued.

In general, if the fill buffer is being used and the next load or store command requires the fill buffer, only one additional command can be accepted before causing additional DCU commands to stall the pipeline.

The DCU can accept up to three outstanding store commands before stalling the CPU pipeline for additional data cache commands.

The DCU can have two flushes pending before stalling the CPU pipeline.

DCU cache operations other than loads and stores stall the CPU pipeline until all prior data cache operations complete. Any subsequent data cache command will stall the pipeline until the prior operation is complete.

4.6.2 Cache Operation Priorities

The DCU uses a priority signal to improve performance when pipeline stalls occur. When the pipeline is stalled because of a data cache operation, the DCU asserts the priority signal to the PLB. The priority signal tells the external bus that the DCU requires immediate service, and is valid only when the data cache is requesting access to the PLB. The priority signal is asserted for all loads that require external data, or when the data cache is requesting the PLB and stalling an operation that is being presented to the data cache.

Table 4-4 provides examples of when the priority is asserted and deasserted.

Table 4-4. Priority Changes With Different Data Cache Operations

Instruction Requesting PLB	Priority	Next Instruction	Priority
Any load from external memory	1	N/A	N/A
Any store	0	Any other cache operation not being accepted by the DCU.	1
dcbf	0	Any cache hit.	0
dcbf/dcbst	0	Load non-cache.	1
dcbf/dcbst	0	Another command that requires a line flush.	1
dcbt	0	Any cache hit.	0
dcbi/dccci/dcbz	0	N/A	N/A

4.6.3 Simultaneous Cache Operations

Some cache operations can occur simultaneously to improve DCU performance. For example, combinations of line fills, line flushes, word load/stores, and operations that hit in the cache can occur simultaneously. Cache operations other than loads/stores cannot begin until the PLB completes all previous operations.

4.6.4 Sequential Cache Operations

Some common cache operations, when performed sequentially, can limit DCU performance: sequential loads/stores to non-cachable storage regions, sequential line fills, and sequential line flushes.

In the case of sequential cache hits, the most commonly occurring operations, the DCU loads or stores data every cycle. In such cases, the DCU does not limit performance.

However, when a load from a non-cachable storage region is followed by multiple loads from non-cachable regions, the loads can complete no faster than every four cycles, assuming that the addresses are accepted during the same cycle in which it is requested, and that the data is returned during the cycle after the load is accepted.

Similarly, when a store to a non-cachable storage region is followed by multiple stores to non-cachable regions the fastest that the stores can complete is every other cycle. The DCU can have accepted up to three stores before additional DCU commands will stall waiting for the prior stores to complete.

Sequential line fills can limit DCU performance. Line fills occur when a load/store or **dcbt** instruction misses in the cache, and can be pipelined on the PLB interface such that up to two requests can be accepted before stalling subsequent requests. The subsequent operations will wait in the DCU until the first line fill completes. The line fills must complete in the order that they are accepted.

Sequential line flushes from the DCU to main memory also limit DCU performance. Flushes occur when a line fill replaces a valid line that is marked dirty (modified), or when a **dcbf** instruction flushes a specific line. If two flushes are pending, the DCU stalls any new data cache operations until the first flush finishes and the second flush begins.

Chapter 5. Fixed-Point Interrupts and Exceptions

An *interrupt* is the action in which the processor saves its old context (MSR and instruction pointer) and begins execution at a pre-determined interrupt-handler address, with a modified MSR.

Exceptions are events which, if enabled, cause the processor to take an interrupt. Exceptions are generated by signals from internal and external peripherals, instructions, internal timer facilities, debug events, or error conditions.

Table 5-2 on page 5-6, lists the interrupts handled by the PPC405 in the order of interrupt vector offsets. Detailed descriptions of each interrupt follow, in the same order. Table 5-2 also provides an index to the descriptions.

Several registers support interrupt handling and control. “General Interrupt Handling Registers” on page 5-7 describes the general interrupt handling registers:

- Data Exception Address Register (DEAR)
- Exception Syndrome Register (ESR)
- Exception Vector Prefix Register (EVPR)
- Machine State Register (MSR)
- Save/Restore Registers (SRR0–SRR3)

Two external interrupt input signals are provided in the PPC405. One external interrupt input is for critical interrupts; the other in for non-critical interrupts. Both external interrupts are maskable. The MSR enables critical and noncritical external interrupt signals.

5.1 Architectural Definitions and Behavior

Precise interrupts are those for which the instruction pointer saved by the interrupt must be either the address of the excepting instruction or the address of the next sequential instruction. *Imprecise* interrupts are those for which it is possible (but not required) for the saved instruction pointer to be something else, possibly prohibiting guaranteed software recovery.

Note that “precise” and “imprecise” are defined assuming that the interrupts are unmasked (enabled to occur) when the associated exception occurs. Consider an exception that would cause a precise interrupt, if the interrupt was enabled at the time of the exception, but that occurs while the interrupt is masked. Some exceptions of this type can cause the interrupt to occur later, immediately upon its enabling. In such a case, the interrupt is not considered precise with respect to the enabling instruction, but imprecise (“delayed precise”) with respect to the cause of the exception.

Asynchronous interrupts are caused by events which are independent of instruction execution. All asynchronous interrupts are precise, and the following rules apply:

1. All instructions prior to the one whose address is reported to the interrupt handling routine (in the save/restore register) have completed execution. However, some storage accesses generated by these preceding instructions may not have completed.
2. No subsequent instruction has begun execution, including the instruction whose address is reported to the interrupt handling routine.

3. The instruction having its address reported to the interrupt handler may appear not to have begun execution, or may have partially completed.

Synchronous interrupts are caused directly by the execution (or attempted execution) of instructions. Synchronous interrupts can be either precise or imprecise.

For synchronous precise interrupts, the following rules apply:

1. The save/restore register addresses either the instruction causing the exception or the next sequential instruction. Which instruction is addressed is determined by the interrupt type and status bits.
2. All instructions preceding the instruction causing the exception have completed execution. However, some storage accesses generated by these preceding instructions may not have completed.
3. The instruction causing the exception may appear not to have begun execution (except for causing the exception), may have partially completed, or may have completed, depending on the interrupt type.
4. No subsequent instruction has begun execution.

Refer to *IBM PowerPC Embedded Environment* for an architectural description of imprecise interrupts.

Machine check interrupts are a special case typically caused by some kind of hardware or storage subsystem failure, or by an attempt to access an invalid address. A machine check can be indirectly caused by the execution of an instruction, but not recognized 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 thought of as synchronous, nor as precise or imprecise. For machine checks, the following general rules apply:

1. No instruction following the one whose address is reported to the machine check handler in the save/restore register has begun execution.
2. The instruction whose address is reported to the machine check handler in the save/restore register, and all previous instructions, may or may not have completed successfully. All previous instructions that would ever complete have completed, within the context existing before the machine check interrupt. No further interrupt (other than possible additional machine checks) can occur as a result of those instructions.

5.2 Behavior of the PPC405 Processor Core Implementation

All interrupts, except for machine checks, are handled precisely. Precise handling implies that the address of the excepting instruction (for synchronous exceptions other than the system call exception), or the address of the next instruction to be executed (asynchronous exceptions and the system call exception), is passed to an interrupt handling routine. Precise handling also implies that all instructions that precede the instruction whose address is reported to the interrupt handling routine have executed and that no subsequent instruction has begun execution. The specific instruction whose address is reported may not have begun execution or may have partially completed, as specified for each precise interrupt type.

Synchronous precise interrupts include most debug event interrupts, program interrupts, instruction and data storage interrupts, auxiliary processor unit (APU) interrupts, floating point unit (FPU) interrupts, TLB miss interrupts, system call interrupts, and alignment interrupts.

Asynchronous precise interrupts include the critical and noncritical external interrupts, timer facility interrupts, and some debug events.

In the PPC405, machine checks are handled as critical interrupts (see “Critical and Noncritical Interrupts” on page 5-5). If a machine check is associated with an instruction fetch, the critical interrupt save/restore register contains the address of the instruction being fetched when the machine check occurred.

The synchronism of instruction-side machine checks (errors that occur while attempting to fetch an instruction from external memory) requires further explanation. Fetch requests to cachable memory that miss in the instruction cache unit (ICU) cause an instruction cache line fill (eight words). If any instructions (words) in the fetched line are associated with an exception, an interrupt occurs upon attempted execution and the cache line is invalidated.

It is improper to declare an exception when an erroneous word is passed to the fetcher; the address could be the result of an incorrect speculative access. It is quite likely that no attempt will be made to execute an instruction from the erroneous address. An instruction-side machine check interrupt occurs only when execution is attempted. If an exception occurs, execution is suppressed, SRR2 contains the erroneous address, and the ESR indicates that an instruction-side machine check occurred. Although such an interrupt is clearly asynchronous to the erroneous memory access, it is handled synchronously with respect to the attempted execution from the erroneous address.

Except for machine checks, all PPC405 interrupts are handled precisely:

- The address of the excepting instruction (for synchronous exceptions, other than the system call exception) or the address of the next sequential instruction (for asynchronous exceptions and the system call exception) is passed to the interrupt handling routine.
- All instructions that precede the instruction whose address is reported to the interrupt handling routine have completed execution and that no subsequent instruction has begun execution. The specific instruction whose address is reported might not have begun execution or might have partially completed, as specified for each interrupt type.

5.3 Interrupt Handling Priorities

The PPC405 core only one interrupt at a time. Multiple simultaneous interrupts are handled in the priority order shown in Table 5-1 (assuming, of course, that the interrupt types are enabled).

Multiple interrupts can exist simultaneously, each of which requires the generation of an interrupt. The architecture does not provide for simultaneously reporting more than one interrupt of the same class (critical or non-critical). Therefore, interrupts are ordered with respect to each other. A masking mechanism is available for certain persistent interrupt types.

When an interrupt type is masked, and an event causes an exception which would normally generate an interrupt of that type, the exception *persists* as a *status* bit in a register. 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 is finally generated.

All asynchronous interrupt types can be masked. In addition, certain synchronous interrupt types can be masked.

Table 5-1. Interrupt Handling Priorities

Priority	Interrupt Type	Critical or Noncritical	Causing Conditions
1	Machine check—data	Critical	External bus error during data-side access
2	Debug—IAC	Critical	IAC debug event (in internal debug mode)
3	Machine check—instruction	Critical	Attempted execution of instruction for which an external bus error occurred during fetch
4	Debug—EXC, UDE	Critical	EXC or UDE debug event (in internal debug mode)
5	Critical interrupt input	Critical	Active level on the critical interrupt input
6	Watchdog timer—first time-out	Critical	Posting of an enabled first time-out of the watchdog timer in the TSR
7	Instruction TLB Miss	Noncritical	Attempted execution of an instruction at an address and process ID for which a valid matching entry was not found in the TLB
8	Instruction storage — ZPR[Zn] = 00	Noncritical	Instruction translation is active, execution access to the translated address is not permitted because ZPR[Zn] = 00 in user mode, and an attempt is made to execute the instruction
9	Instruction storage — TLB_entry[EX] = 0	Noncritical	Instruction translation is active, execution access to the translated address is not permitted because TLB_entry[EX] = 0, and an attempt is made to execute the instruction
	Instruction storage — TLB_entry[G] = 1 or SGR[Gn] = 1	Noncritical	Instruction translation is active, the page is marked guarded, and an attempt is made to execute the instruction
	Program	Noncritical	Attempted execution of illegal instructions, TRAP instruction, privileged instruction in problem state, or auxiliary processor (APU) instruction, or unimplemented FPU instruction, or unimplemented APU instruction, or APU interrupt, or FPU interrupt
	System call	Noncritical	Execution of the sc instruction
	APU Unavailable	Noncritical	Attempted execution of an APU instruction when MSR[AP] = 0
	FPU Unavailable	Noncritical	Attempted execution of an FPU instruction when MSR[FP]=0.
11	Data TLB miss	Noncritical	Valid matching entry for the effective address and process ID of an attempted data access is not found in the TLB
12	Data storage— ZPR[Zn] = 00	Noncritical	Data translation is active and data-side access to the translated address is not permitted because ZPR[Zn] = 00 in user mode

Table 5-1. Interrupt Handling Priorities (continued)

Priority	Interrupt Type	Critical or Noncritical	Causing Conditions
13	Data storage— TLB_entry[WR] = 0	Noncritical	Data translation is active and write access to the translated address is not permitted because TLB_entry[WR] = 0
	Data storage— TLB_entry[U0] = 1 or SU0R[U n] = 1	Noncritical	Data translation is active and write access to the translated address is not permitted because TLB_entry[U0] = 1 or SU0R[U n] = 1
14	Alignment	Noncritical	dcbz to non-cachable address or write-through storage; non-word aligned dcread , lwarx , and stwcx , as described in Table 5-10; misaligned APU or FPU data access
15	Debug—BT, DAC, DVC, IC, TIE	Critical	BT, DAC, DVC, IC, TIE debug event (in internal debug mode)
16	External interrupt input	Noncritical	Interrupts from the external interrupt input
17	Fixed Interval Timer (FIT)	Noncritical	Posting of an enabled FIT interrupt in the TSR
18	Programmable Interval Timer (PIT)	Noncritical	Posting of an enabled PIT interrupt in the TSR

5.4 Critical and Noncritical Interrupts

The PPC405 processes interrupts as noncritical and critical. The following interrupts are defined as *noncritical*: data storage, instruction storage, an active external interrupt input, alignment, program, FPU unavailable, APU unavailable, system call, programmable interval timer (PIT), fixed interval timer (FIT), data TLB miss, and instruction TLB miss. The following interrupts are defined as *critical*: machine check interrupts (instruction- and data-side), debug interrupts, interrupts caused by an active critical interrupt input, and the first time-out from the watchdog timer.

When a *noncritical* interrupt is taken, Save/Restore Register 0 (SRR0) is written with the address of the excepting instruction (most synchronous interrupts) or the next sequential instruction to be processed (asynchronous interrupts and system call).

If the PPC405 was executing a multicycle instruction (multiply, divide, or cache operation), the instruction is terminated and its address is written in SRR0.

Aligned scalar loads/stores that are interrupted do not appear on the PLB. An aligned scalar load/store cannot be interrupted after it is requested on the PLB, so the Guarded (G) storage attribute does not need to prevent the interruption of an aligned scalar load/store.

To enhance performance, the DCU can respond to non-cachable load requests by retrieving a line instead of a word. This is controlled by CCR0[LWL]. Note, however, that if CCR0[LWL] = 1, and the target non-cachable region is also marked as guarded (the G storage attribute is set to 1), that the DCU will request on the PLB only those bytes requested by the CPU.

Load/store multiples, load/store string, and misaligned scalar loads/stores that cross a word boundary can be interrupted and restarted upon return from the interrupt handler.

When load instructions terminate, the addressing registers are not updated. This ensures that the instructions can be restarted; if the addressing registers were in the range of registers to be loaded, this would be an invalid form in any event. Some target registers of a load instruction may have been

written by the time of the interrupt; when the instruction restarts, the registers will simply be written again. Similarly, some of the target memory of a store instruction may have been written, and is written again when the instruction restarts.

Save/Restore Register 1 (SRR1) is written with the contents of the MSR; the MSR is then updated to reflect the new machine context. The new MSR contents take effect beginning with the first instruction of the interrupt handling routine.

Interrupt handling routine instructions are fetched at an address determined by the interrupt type. The address of the interrupt handling routine is formed by concatenating the 16 high-order bits of the EVPR and the interrupt vector offset. (A user must initialize the EVPR contents at power-up using an **mtspr** instruction.)

Table 5-2 shows the interrupt vector offsets for the interrupt types. Note that there can be multiple sources of the same interrupt type; interrupts of the same type are mapped to the same interrupt vector, regardless of source. In such cases, the interrupt handling routine must examine status registers to determine the exact source of the interrupt.

At the end of the interrupt handling routine, execution of an **rfi** instruction forces the contents of SRR0 and SRR1 to be written to the program counter and the MSR, respectively. Execution then begins at the address in the program counter.

Critical interrupts are processed similarly. When a critical interrupt is taken, Save/Restore Register 2 (SRR2) and Save/Restore Register 3 (SRR3) hold the next sequential address to be processed when returning from the interrupt, and the contents of the MSR, respectively. At the end of the critical interrupt handling routine, execution of an **rfci** instruction writes the contents of SRR2 and SRR3 into the program counter and the MSR, respectively.

Table 5-2. Interrupt Vector Offsets

Offset	Interrupt Type	Interrupt Class	Category	Page
0x0100	Critical input interrupt	Asynchronous precise	Critical	5-13
0x0200	Machine check—data	—	Critical	5-14
	Machine check—instruction	—	Critical	5-14
0x0300	Data storage interrupt— MSR[DR]=1 and ZPR[Zn] = 0 or TLB_entry[WR] = 0 or TLB_entry[U0] = 1 or SU0R[Un] = 1	Synchronous precise	Noncritical	5-16
0x0400	Instruction storage interrupt	Synchronous precise	Noncritical	5-17
0x0500	External interrupt	Asynchronous precise	Noncritical	5-18
0x0600	Alignment	Synchronous precise	Noncritical	5-19
0x0700	Program	Synchronous precise	Noncritical	5-20
0x0800	FPU Unavailable	Synchronous precise	Noncritical	5-21
0x0C00	System Call	Synchronous precise	Noncritical	5-22
0x0F20	APU Unavailable	Synchronous precise	Noncritical	5-22
0x1000	PIT	Asynchronous precise	Noncritical	5-22
0x1010	FIT	Asynchronous precise	Noncritical	5-23
0x1020	Watchdog timer	Asynchronous precise	Critical	5-24
0x1100	Data TLB miss	Synchronous precise	Noncritical	5-25

Table 5-2. Interrupt Vector Offsets (continued)

Offset	Interrupt Type	Interrupt Class	Category	Page
0x1200	Instruction TLB miss	Synchronous precise	Noncritical	5-25
0x2000	Debug—BT, DAC, DVC, IAC, IC, TIE	Synchronous precise	Critical	5-26
	Debug—EXC, UDE	Asynchronous precise	Critical	

5.5 General Interrupt Handling Registers

The general interrupt handling registers are the Machine State Register (MSR), SRR0–SRR3, the Exception Vector Prefix Register (EVPR), the Exception Syndrome Register (ESR), and the Data Exception Address Register (DEAR).

5.5.1 Machine State Register (MSR)

The MSR is a 32-bit register that holds the current context of the PPC405. When a noncritical interrupt is taken, the MSR contents are written to SRR1; when a critical interrupt is taken, the MSR contents are written to SRR3. When an **rfi** or **rfci** instruction executes, the contents of the MSR are read from SRR1 or SRR3, respectively.

Programming Note: The **rfi** and **rfci** instructions can alter reserved MSR fields.

The MSR contents can be read into a general purpose register (GPRs) using an **mfmsr** instruction. The contents of a GPR can be written to the MSR using an **mtmsr** instruction. The MSR[EE] bit may be set/cleared atomically using the **wrtee** or **wrteei** instructions.

Figure 5-1 shows the MSR bit definitions and describes the function of each bit.




Figure 5-1. Machine State Register (MSR)

0:5		Reserved	
6	AP	Auxiliary Processor Available 0 APU not available. 1 APU available.	
7:11		Reserved	
12	APE	APU Exception Enable 0 APU exception disabled. 1 APU exception enabled.	
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 interrupts are disabled. 1 Critical interrupts are enabled.	Controls the critical interrupt input and watchdog timer first time-out interrupts.
15		Reserved	
16	EE	External Interrupt Enable 0 Asynchronous interrupts are disabled. 1 Asynchronous interrupts are enabled.	Controls the non-critical external interrupt input, PIT, and FIT interrupts.
17	PR	Problem State 0 Supervisor state (all instructions allowed). 1 Problem state (some instructions not allowed).	
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 Debug wait mode is disabled. 1 Debug wait mode is enabled.	
22	DE	Debug Interrupts 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	IR	Instruction Relocate 0 Instruction address translation is disabled. 1 Instruction address translation is enabled.	
27	DR	Data Relocate 0 Data address translation is disabled. 1 Data address translation is enabled.	

28:31	Reserved
-------	----------

5.5.2 Save/Restore Registers 0 and 1 (SRR0–SRR1)

SRR0 and SRR1 are 32-bit registers that hold the interrupted machine context when a noncritical interrupt is processed. On interrupt, SRR0 is set to the current or next instruction address and the contents of the MSR are written to SRR1. When an **rfi** instruction is executed at the end of the interrupt handler, the program counter and the MSR are restored from SRR0 and SRR1, respectively.

The contents of SRR0 and SRR1 can be written into GPRs using the **mf spr** instruction. The contents of GPRs can be written to SRR0 and SRR1 using the **mt spr** instruction.

Figure 5-2 shows the bit definitions for SRR0.



Figure 5-2. Save/Restore Register 0 (SRR0)

0:29	SRR0 receives an instruction address when a non-critical interrupt is taken; the Program Counter is restored from SRR0 when rfi executes.
30:31	Reserved

Figure 5-3 shows the bit definitions for SRR1.




Figure 5-3. Save/Restore Register 1 (SRR1)

0:31	SRR1 receives a copy of the MSR when an interrupt is taken; the MSR is restored from SRR1 when rfi executes.
------	---

5.5.3 Save/Restore Registers 2 and 3 (SRR2–SRR3)

SRR2 and SRR3 are 32-bit registers that hold the interrupted machine context when a critical interrupt is processed. On interrupt, SRR2 is set to the current or next instruction address and the contents of the MSR are written to SRR3. When an **rfci** instruction is executed at the end of the interrupt handler, the program counter and the MSR are restored from SRR2 and SRR3, respectively.

The contents of SRR2 and SRR3 can be written to GPRs using the **mf spr** instruction. The contents of GPRs can be written to SRR2 and SRR3 using the **mt spr** instruction.

Figure 5-4 shows the bit definitions for SRR2.



Figure 5-4. Save/Restore Register 2 (SRR2)

0:29	SRR2 receives an instruction address when a critical interrupt is taken; the Program Counter is restored from SRR2 when r fci executes.
30:31	Reserved

Figure 5-5 shows the bit definitions for SRR3.




Figure 5-5. Save/Restore Register 3 (SRR3)

0:31	SRR3 receives a copy of the MSR when a critical interrupt is taken; the MSR is restored from SRR3 when r fci executes.
------	---

Because critical interrupts do not automatically clear MSR[ME], SRR2 and SRR3 can be corrupted by a machine check interrupt, if the machine check occurs while SRR2 and SRR3 contain valid data that has not yet been saved by the critical interrupt handler.

5.5.4 Exception Vector Prefix Register (EVPR)

The EVPR is a 32-bit register whose high-order 16 bits contain the prefix for the address of an interrupt handling routine. The 16-bit interrupt vector offsets (shown in Table 5-2 on page 5-6) are concatenated to the right of the high-order 16 bits of the EVPR to form the 32-bit address of an interrupt handling routine.

The contents of the EVPR can be written to a GPR using the **mf spr** instruction. The contents of a GPR can be written to EVPR using the **mt spr** instruction.

Figure 5-6 shows the EVPR bit definitions.




Figure 5-6. Exception Vector Prefix Register (EVPR)

0:15	EVP	Exception Vector Prefix
16:31		Reserved

5.5.5 Exception Syndrome Register (ESR)

The ESR is a 32-bit register whose bits help to specify the exact cause of various synchronous interrupts. These interrupts include instruction side machine checks, data storage interrupts, and program interrupts, instruction storage interrupts, and data TLB miss interrupts.

“Instruction Machine Check Handling” on page 5-14 describes instruction machine checks. “Data Storage Interrupt” on page 5-16 describes data storage interrupts. “Program Interrupt” on page 5-20 describes program interrupts.

Although interrupt handling routines are not required to reset the ESR, it is recommended that instruction machine check handlers reset the ESR; “Instruction Machine Check Handling” on page 5-14 describes why such resets are recommended.

The contents of the ESR can be written to a GPR using the **mf spr** instruction. The contents of a GPR can be written to the ESR using the **mt spr** instruction.

Figure 5-7 shows the ESR bit definitions.




Figure 5-7. Exception Syndrome Register (ESR)

0	MCI	Machine check—instruction 0 Instruction machine check did not occur. 1 Instruction machine check occurred.
1:3		Reserved
4	PIL	Program interrupt—illegal 0 Illegal Instruction error did not occur. 1 Illegal Instruction error occurred.
5	PPR	Program interrupt—privileged 0 Privileged instruction error did not occur. 1 Privileged instruction error occurred.

6	PTR	Program interrupt—trap 0 Trap with successful compare did not occur. 1 Trap with successful compare occurred.
7	PEU	Program interrupt—Unimplemented 0 APU/FPU unimplemented exception did not occur. 1 APU/FPU unimplemented exception occurred.
8	DST	Data storage interrupt—store fault 0 Excepting instruction was not a store. 1 Excepting instruction was a store (includes dcbi , dcbz , and dccci).
9	DIZ	Data/instruction storage interrupt—zone fault 0 Excepting condition was not a zone fault. 1 Excepting condition was a zone fault.
10:11		Reserved
12	PFP	Program interrupt—FPU 0 FPU interrupt did not occur. 1 FPU interrupt occurred.
13	PAP	Program interrupt—APU 0 APU interrupt did not occur. 1 APU interrupt occurred.
14:15		Reserved
16	U0F	Data storage interrupt—U0 fault 0 Excepting instruction did not cause a U0 fault. 1 Excepting instruction did cause a U0 fault.
17:31		Reserved

In general, ESR bits are set to indicate the type of precise interrupt that occurred; other bits are cleared. However, the machine check—instruction (ESR[MCI]) bit behaves differently. Because instruction-side machine checks can occur without an interrupt being taken (if MSR[ME] = 0), ESR[MCI] can be set even while other ESR-setting interrupts (program, data storage, DTLB-miss) occurring. Thus, data storage and program interrupts leave ESR[MCI] unchanged, clear all other ESR bits, and set the bits associated with any data storage or program interrupts that occurred. Enabled instruction-side machine checks (MSR[ME] = 1) set ESR[MCI] and clear the data storage and program interrupt bits.

If a machine check—instruction interrupt occurs but is disabled (MSR[ME] = 0), it sets ESR[MCI] but leaves the data storage and program interrupt bits alone. If a machine check—instruction interrupt occurs while MSR[ME] = 0, *and* the instruction upon which the machine check—instruction interrupt is occurring also is some other kind of ESR-setting instruction (program, data storage, DTLB-miss, or instruction storage interrupt), ESR[MCI] is set to indicate that a machine check—instruction interrupt

occurred; the other ESR bits are set or cleared to indicate the other interrupt. These scenarios are summarized in Table 5-3

Table 5-3. ESR Alteration by Various Interrupts

Scenario	ECR[MCI]	ESR _{4:7, 12:13}	ESR _{8:9, 16}
Program interrupt	Unchanged	Set to type	Cleared
Data storage interrupt	Unchanged	Cleared	Set to Type
Data TLB miss interrupt	Unchanged	Cleared	Cleared
Machine check—instruction	Set to 1	Cleared	Cleared
Disabled MCI, no others	Unchanged	Unchanged	Unchanged
Disabled MCI and program interrupt	Unchanged	Set to type	Cleared

Engineering Note: An implementation can use additional ESR bits to identify implementation-specific exception types. Implementations can also use the ESR to record information about the cause of a machine check interrupt.

5.5.6 Data Exception Address Register (DEAR)

The DEAR is a 32-bit register that contains the address of the access for which one of the following synchronous precise errors occurred: alignment error, data TLB miss, or data storage interrupt.

The contents of the DEAR can be written to a GPR using the **mf spr** instruction. The contents of a GPR can be written to the DEAR using the **mt spr** instruction.

Figure 5-8 shows the DEAR bit definitions.



Figure 5-8. Data Exception Address Register (DEAR)

0:31		Address of Data Error (synchronous)
------	--	-------------------------------------

5.6 Critical Input Interrupts

An external source requests a critical interrupt by driving the critical interrupt input active. The critical interrupt is recognized if enabled by MSR[CE].

MSR[CE] also enables the watchdog timer first-time-out interrupt. However, the watchdog interrupt has a different interrupt vector than the critical pin interrupt. See “Watchdog Timer Interrupt” on page 5-24.

After detecting a critical interrupt, if no synchronous precise interrupts are outstanding, the PPC405 immediately takes the critical interrupt and writes the address of the next instruction to be executed in

SRR2. Simultaneously, the contents of the MSR are saved in SRR3. MSR[CE] is reset to 0 to prevent another critical interrupt or the watchdog timer first time-out interrupt from interrupting the critical interrupt handler before SRR2 and SRR3 get saved. MSR[DE] is reset to 0 to disable debug interrupts during the critical interrupt handler.

The MSR is also written with the values shown in Table 5-4 on page 5-14. The high-order 16 bits of the program counter are then loaded with the contents of the EVPR and the low-order 16 bits of the program counter are loaded with 0x0100. Interrupt processing begins at the address in the program counter.

Inside the interrupt handling routine, after the contents of SRR2/SRR3 are saved, critical interrupts can be enabled again by setting MSR[CE] = 1.

Executing an **rfci** instruction restores the program counter from SRR2 and the MSR from SRR3, and execution resumes at the address in the program counter.

Table 5-4. Register Settings during Critical Input Interrupts

SRR2	Written with the address of the next instruction to be executed
SRR3	Written with the contents of the MSR
MSR	AP, APE, WE, CE, EE, PR, FP, FE0, DWE, DE, FE1, IR, DR←0 ME← unchanged
PC	EVPR[0:15] 0x0100

5.7 Machine Check Interrupts

When an external bus error occurs on an instruction fetch, and execution of that instruction is subsequently attempted, a machine check—instruction interrupt occurs.

When an external bus error occurs while attempting data accesses, a machine check—data interrupt occurs.

When an instruction-side machine check interrupt occurs, the PPC405 stores the address of the excepting instruction in SRR2. When a data-side machine check occurs, the PPC405 stores the address of the next sequential instruction in SRR2. Simultaneously, for all machine check interrupts, the contents of the MSR are loaded into SRR3.

The MSR Machine Check Enable bit (MSR[ME]) is reset to 0 to disable another machine check from interrupting the machine check interrupt handling routine. The other MSR bits are loaded with the values shown in Table 5-5 on page 5-15 and Table 5-6 on page 5-15. The high-order 16 bits of the program counter are then written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x0200. Interrupt processing begins at the new address in the program counter.

Executing an **rfci** instruction restores the program counter from SRR2 and the MSR from SRR3, and execution resumes at the address in the program counter.

5.7.1 Instruction Machine Check Handling

When a machine check occurs on an instruction fetch, *and execution of that instruction is subsequently attempted*, a machine check—instruction interrupt occurs. If enabled by MSR[ME], the processor reports the machine check—instruction interrupt by vectoring to the machine check

handler (EVPR[0:15] || 0x0200), setting ESR[MCI]. Note that only a bus error can cause a machine check—instruction interrupt. Taking the vector automatically clears MSR[ME] and the other MSR fields.

Note that it is improper to declare a machine check—instruction interrupt when the instruction is fetched, because the address is possibly the result of an incorrect speculation by the fetcher. It is quite likely that no attempt will be made to execute an instruction from the erroneous address. The interrupt will occur only if execution of the instruction is subsequently attempted.

When a machine check occurs on an instruction fetch, the erroneous instruction is never validated in the instruction cache unit (ICU). Fetch requests to cachable memory that miss in the ICU cause an instruction cache line fill (eight words). If any words in the fetched line are associated with an error, an interrupt occurs upon attempted execution and the cache line is invalidated. If any word in the line is in error, the cache line is invalidated after the line fill.

ESR[MCI] is set, even if MSR[ME] = 0. This means that if a machine check—instruction interrupt occurs while running in code in which MSR[ME] is disabled, the machine check—instruction interrupt is recorded in the ESR, but no interrupt occurs. Software running with MSR[ME] disabled can sample ESR[MCI] to determine whether at least one machine check—instruction interrupt occurred during the disabled execution.

If a new machine check—instruction interrupt occurs after MSR[ME] is enabled again, the new machine check—instruction interrupt is recorded in ESR[MCI] and the machine check—instruction interrupt handler is invoked. However, enabling MSR[ME] again does *not* cause a machine Check interrupt to occur simply due to the presence of ESR[MCI] indicating that a machine check—instruction interrupt occurred while MSR[ME] was disabled. The machine check—instruction interrupt must occur while MSR[ME] is enabled for the machine check interrupt to be taken. Software should, in general, clear the ESR bits before returning from a machine check interrupt to avoid any ambiguity when handling subsequent machine check interrupts.

Table 5-5. Register Settings during Machine Check—Instruction Interrupts

SRR2	Written with the address that caused the machine check.
SRR3	Written with the contents of the MSR
MSR	WE, CE, EE, PR, ME, FP, FE0, DWE, DE, FE1, IR, DR←0
PC	EVPR[0:15] 0x0200
ESR	MCI ← 1 All other bits are cleared.

5.7.2 Data Machine Check Handling

When a machine check occurs on a data access, a machine check—data interrupt occurs. The handling of machine check—data interrupts is implementation-specific.

Table 5-6. Register Settings during Machine Check—Data Interrupts

SRR2	Written with the address of the next sequential instruction.
SRR3	Written with the contents of the MSR
MSR	WE, CE, EE, PR, ME, FP, FE0, DWE, DE, FE1, IR, DR←0
PC	EVPR[0:15] 0x0200

5.8 Data Storage Interrupt

The data storage interrupt occurs when the desired access to the effective address is not permitted for any of the following reasons:

- A U0 fault: any store to an EA with the U0 storage attribute set and CCR0[U0XE] = 1
- In the problem state with data translation enabled:
 - A *zone fault*, which is any user-mode storage access (data load, store, **icbi**, **dcbz**, **dcbst**, or **dcbf**) with an effective address with (ZPR field) = 00. (**dcbt** and **dcbtst** will no-op in this situation, rather than cause an interrupt. The instructions **dcbi**, **dccci**, **icbt**, and **iccci**, being privileged, cannot cause zone fault data storage interrupts.)
 - Data store or **dcbz** to an effective address with the WR bit clear and (ZPR field) ≠ 11. (The privileged instructions **dcbi** and **dccci** are treated as “stores,” but will cause privileged program interrupts, rather than data storage interrupts.)
- In the supervisor state with data translation enabled:
 - Data store, **dcbi**, **dcbz**, or **dccci** to an effective address with the WR bit clear and (ZPR field) other than 11 or 10.

Programming Note: The **icbi**, **icbt**, and **iccci** instructions are treated as loads from the addressed byte with respect to address translation and protection. Instruction cache operations use MSR[DR], not MSR[IR], to determine translation of their operands. Instruction storage interrupts and Instruction-side TLB Miss Interrupts are associated with the *fetching* of instructions, not with the execution of instructions. Data storage interrupts and data TLB miss interrupts are associated with the *execution* of instruction cache operations.

When a data storage interrupt is detected, the PPC405 suppresses the instruction causing the interrupt and writes the instruction address in SRR0. The Data Exception Address Register (DEAR) is loaded with the data address that caused the access violation. ESR bits are loaded as shown in Table 5-7 on page 5-17 to provide further information about the error. The current contents of the MSR are loaded into SRR1, and MSR bits are then loaded with the values shown in Table 5-7.

The high-order 16 bits of the program counter are then loaded with the contents of the EVPR and the low-order 16 bits of the program counter are loaded with 0x0300. Interrupt processing begins at the new address in the program counter. Executing the return from interrupt instruction (**rfi**) restores the contents of the program counter and the MSR from SRR0 and SRR1, respectively, and the PPC405 resumes execution at the new program counter address.

For instructions that can simultaneously generate program interrupts (privileged instructions executed in Problem State) and data storage interrupts, the program interrupt has priority.

Table 5-7. Register Settings during Data Storage Interrupts

SRR0	Written with the EA of the instruction causing the data storage interrupt
SRR1	Written with the value of the MSR at the time of the interrupt
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR←0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0300
DEAR	Written with the EA of the failed access
ESR	DST ← 1 if excepting operation is a store DIZ ← 1 if access failure caused by a zone protection fault (ZPR[Zn] = 00 in user mode) UOF ← 1 if access failure caused by a U0 fault (the U0 storage attribute is set and CCR0[U0XE] = 1) MCI ← unchanged All other bits are cleared.

5.9 Instruction Storage Interrupt

The instruction storage interrupt is generated when instruction translation is active and execution is attempted for an instruction whose fetch access to the effective address is not permitted for any of the following reasons:

- In Problem State:
 - Instruction fetch from an effective address with (ZPR field) = 00.
 - Instruction fetch from an effective address with the EX bit clear and (ZPR field) ≠ 11.
 - Instruction fetch from an effective address contained within a Guarded region (G=1).
- In Supervisor State:
 - Instruction fetch from an effective address with the EX bit clear and (ZPR field) other than 11 or 10.
 - Instruction fetch from an effective address contained within a Guarded region (G=1).

SRR0 will save the address of the instruction causing the instruction storage interrupt.

ESR is set to indicate the following conditions:

- If ESR[DIZ] = 1, the excepting condition was a zone fault: the attempted execution of an instruction address fetched in user-mode with (ZPR field) = 00.
- If ESR[DIZ] = 0, then the excepting condition was either EX = 0 or G = 1.

The interrupt is precise with respect to the attempted execution of the instruction. Program flow vectors to EVPR[0:15] || 0x0400.

The following registers are modified to the specified values:

Table 5-8. Register Settings during Instruction Storage Interrupts

SRR0	Set to the EA of the instruction for which execute access was not permitted
SRR1	Set to the value of the MSR at the time of the interrupt
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0400
ESR	DIZ ← 1 if access failure due to a zone protection fault (ZPR[Zn] = 00 in user mode) Note: If ESR[DIZ] is not set, the interrupt occurred because TBL_entry[EX] was clear in an otherwise accessible zone, or because of an instruction fetch from a storage region marked as guarded. See “Exception Syndrome Register (ESR)” on page 5-11 for details of ESR operation. MCI ← unchanged All other bits are cleared.

5.10 External Interrupt

External interrupts are triggered by active levels on the external interrupt inputs. All external interrupting events are presented to the processor as a single external interrupt. External interrupts are enabled or disabled by MSR[EE].

Programming Note: MSR[EE] also enables PIT and FIT interrupts. However, after timer interrupts, control passes to different interrupt vectors than for the interrupts discussed in the preceding paragraph. Therefore, these timer interrupts are described in “Programmable Interval Timer (PIT) Interrupt” on page 5-22 and “Fixed Interval Timer (FIT) Interrupt” on page 5-23.

5.10.1 External Interrupt Handling

When MSR[EE] = 1 (external interrupts are enabled), a noncritical external interrupt occurs, and this interrupt is the highest priority interrupt condition, the processor immediately writes the address of the next sequential instruction into SRR0. Simultaneously, the contents of the MSR are saved in SRR1.

When the processor takes a noncritical external interrupt, MSR[EE] is set to 0. This disables other external interrupts from interrupting the interrupt handler before SRR0 and SRR1 are saved. The MSR is also written with the other values shown in Table 5-9 on page 5-19. The high-order 16 bits of the program counter are written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x0500. Interrupt processing begins at the address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-9. Register Settings during External Interrupts

SRR0	Written with the address of the next sequential instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0500

5.11 Alignment Interrupt

Alignment interrupts are caused by **dcbz** instructions to non-cachable or write-through storage, misaligned **dcread**, **lwarx**, or **stwx**. instructions, or misaligned APU or FPU loads/stores. Table 5-10 summarizes the instructions and conditions causing alignment interrupts.

Table 5-10. Alignment Interrupt Summary

Instructions Causing Alignment Interrupts	Conditions
dcbz	EA in non-cachable or write-through storage
dcread, lwarx, stwcx.	EA not word-aligned
APU or FPU load/store halfword	EA not halfword-aligned
APU or FPU load/store word	EA not word-aligned
APU or FPU load/store doubleword	EA not word-aligned
APU load/store quadword	EA not quadword-aligned

Execution of an instruction causing an alignment interrupt is prohibited from completing. SRR0 is written with the address of that instruction and the current contents of the MSR are saved into SRR1. The DEAR is written with the address that caused the alignment error. The MSR bits are written with the values shown in Table 5-11 on page 5-19. The high-order 16 bits of the program counter are written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x0600. Interrupt processing begins at the new address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter

Alignment interrupts cannot be disabled. To avoid overwrites of SRR0 and SRR1 by alignment interrupts that occur within a handler, interrupt handlers should save these registers as soon as possible.

Table 5-11. Register Settings during Alignment Interrupts

SRR0	Written with the address of the instruction causing the alignment interrupt
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged

Table 5-11. Register Settings during Alignment Interrupts (continued)

PC	EVPR[0:15] 0x0600
DEAR	Written with the address that caused the alignment violation

5.12 Program Interrupt

Program interrupts are caused by attempting to execute:

- An illegal instruction
- A privileged instruction while in the problem state
- Executing a trap instruction with conditions satisfied
- An unimplemented APU or FPU instruction
- An APU instruction with APU interrupt enabled
- An FPU instruction with FPU interrupt enabled

The ESR bits that differentiate these situations are listed and described in Table 5-12. When a program interrupt occurs, the appropriate bit is set and the others are cleared. These interrupts are not maskable.

Table 5-12. ESR Usage for Program Interrupts

Bits	Interrupts	Cause
ESR[PIL]	Illegal instruction	Opcode not recognized
ESR[PPR]	Privileged instruction	Attempt to use a privileged instruction in the problem state
ESR[PTR]	Trap	Excepting instruction is a trap
ESR[PEU]	Unimplemented	An FPU or APU instruction is unimplemented
ESR[PFP]	FPU	Excepting instruction is an FPU instruction
ESR[PAP]	APU	Excepting instruction is an APU instruction

The program interrupt handler does not need to reset the ESR.

When one of the following occurs, the PPC405 does not execute the instruction, but writes the address of the excepting instruction into SRR0:

- Attempted execution of a privileged instruction in problem state
- Attempted execution of an illegal instruction (including memory management instructions when memory management is disabled or when TIEc405MmuEn = 0).

When the TIEc405MmuEn signal is tied to 0, the TLB instructions (**tlbia**, **tlbre**, **tlbsx**, **tlbsync**, and **tlbwe**) are treated as illegal instructions. When execution of any of these instructions occurs under this circumstance, a program interrupt results. Trap instructions can be used as a program interrupt or a debug event, or both (see “Debug Events” on page 8-10 for information about debug events). When a trap instruction is detected as a program interrupt, the PPC405 writes the address of the trap instruction into SRR0. See **tw** on page 9-190 and **twi** on page 9-193 (both in Chapter 9, “Instruction Set”) for a detailed discussion of the behavior of trap instructions with various interrupts enabled.

Attempted execution of an APU instruction while the APUC405exception signal is asserted) results in a program interrupt. Similarly, attempted execution of an FPU instruction while the FPUc405exception signal is asserted) also results in a program interrupt. The following also result in program interrupts: attempted execution of an APU instruction while APUC405DcdAPUOp is asserted but APUC405DcdValidOp is deasserted; and attempted execution of an FPU instruction while APUC405DcdFpuOp but APUC405DcdValidOp is deasserted.

After any program interrupt, the contents of the MSR are $MSR[APA] = 0$, an attempt to execute an instruction intended for an APU causes a program interrupt if $MSR[APE] = 0$ is written into SRR1 and the MSR bits are written with the values shown in Table 5-13. The high-order 16 bits of the program counter are written with the contents of the EVPR; the low-order 16 bits of the program counter are written with 0x0700. Interrupt processing begins at the new address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-13. Register Settings during Program Interrupts

SRR0	Written with the address of the excepting instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0700
ESR	Written with the type of program interrupt. (See Table 5-12) MCI ← unchanged All other bits are cleared.

5.13 FPU Unavailable Interrupt

If $MSR[FP] = 0$, an attempt to execute an FPU instruction for which an FPU asserts APUC405DcdFpuOp causes an FPU unavailable interrupt. The PPC405 FPU does not execute the instruction, but writes the address of the FPU instruction into SRR0.

After an FPU unavailable interrupt occurs, the contents of the MSR are written into SRR1 and the MSR bits are written with the values shown in Table 5-13. The high-order 16 bits of the program counter are written with the contents of the EVPR; the low-order 16 bits of the program counter are written with 0x0800. Interrupt processing begins at the new address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-14. Register Settings during FPU Unavailable Interrupts

SRR0	Written with the address of the excepting instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0800

5.14 System Call Interrupt

System call interrupts occur when a **sc** instruction is executed. The PPC405 writes the address of the instruction following the **sc** into SRR0. The contents of the MSR are written into SRR1 and the MSR bits are written with the values shown in Table 5-15. The high-order 16 bits of the program counter are then written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x0C00. Interrupt processing begins at the new address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-15. Register Settings during System Call Interrupts

SRR0	Written with the address of the instruction following the sc instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0C00

5.15 APU Unavailable Interrupt

If MSR[AP] = 0, an attempt to execute an APU instruction for which an APU asserts APU_C405DcdApuOp causes an APU unavailable interrupt. The PPC405 does not execute the instruction, but writes the address of the APU instruction into SRR0.

After an APU unavailable interrupt, the contents of the MSR are written into SRR1 and the MSR bits are written with the values shown in Table 5-16. The high-order 16 bits of the program counter are written with the contents of the EVPR; the low-order 16 bits of the program counter are written with 0x0F20. Interrupt processing begins at the new address in the program counter.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-16. Register Settings during APU Unavailable Interrupts

SRR0	Written with the address of the excepting instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x0F20

5.16 Programmable Interval Timer (PIT) Interrupt

For a discussion of the PPC405 timer facilities, see Chapter 6, “Timer Facilities.” The PIT is described in “Programmable Interval Timer (PIT)” on page 6-4.

If the PIT interrupt is enabled by TCR[PIE] and MSR[EE], the PPC405 initiates a PIT interrupt after detecting a time-out from the PIT. Time-out is detected when, at the beginning of a clock cycle, TSR[PIS] = 1. (This occurs on the cycle after the PIT decrements on a PIT count of 1.) The PPC405 immediately takes the interrupt. The address of the next sequential instruction is saved in SRR0;

simultaneously, the contents of the MSR are written into SRR1 and the MSR is written with the values shown in Table 5-17. The high-order 16 bits of the program counter are then written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x1000. Interrupt processing begins at the address in the program counter.

To clear a PIT interrupt, the interrupt handling routine must clear the PIT interrupt bit, TSR[PIS]. Clearing is performed by writing a word to TSR, using an **mtspr** instruction, that has 1 in bit positions to be cleared and 0 in all other bit positions. The data written to the TSR is not direct data, but a mask; a 1 clears the bit and 0 has no effect.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-17. Register Settings during Programmable Interval Timer Interrupts

SRR0	Written with the address of the next instruction to be executed
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x1000
TSR	PIS ← 1

5.17 Fixed Interval Timer (FIT) Interrupt

For a discussion of the PPC405 timer facilities, see Chapter 6, “Timer Facilities.” The FIT is described in “Fixed Interval Timer (FIT) Interrupt” on page 5-23.

If the FIT interrupt is enabled by TCR[FIE] and MSR[EE], the PPC405 initiates a FIT interrupt after detecting a time-out from the FIT. Time-out is detected when, at the beginning of a clock cycle, TSR[FIS] = 1. (This occurs on the second cycle after the 0 → 1 transition of the appropriate time-base bit.) The PPC405 immediately takes the interrupt. The address of the next sequential instruction is written into SRR0; simultaneously, the contents of the MSR are written into SRR1 and the MSR is written with the values shown in Table 5-18. The high-order 16 bits of the program counter are then written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x1010. Interrupt processing begins at the address in the program counter.

To clear a FIT interrupt, the interrupt handling routine must clear the FIT interrupt bit, TSR[FIS]. Clearing is performed by writing a word to TSR, using an **mtspr** instruction, that has 1 in any bit positions to be cleared and 0 in all other bit positions. The data written to the TSR is not direct data, but a mask; a 1 clears a bit and 0 has no effect.

Executing an **rfi** instruction restores the program counter from SRR0 and the MSR from SRR1, and execution resumes at the address in the program counter.

Table 5-18. Register Settings during Fixed Interval Timer Interrupts

SRR0	Written with the address of the next sequential instruction
SRR1	Written with the contents of the MSR
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
MSR	WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x1010
TSR	FIS ← 1

5.18 Watchdog Timer Interrupt

For a general description of the PPC405 timer facilities, see Chapter 6, “Timer Facilities.” The watchdog timer (WDT) is described in “Watchdog Timer” on page 6-6.

If the WDT interrupt is enabled by TCR[WIE] and MSR[CE], the PPC405 initiates a WDT interrupt after detecting the first WDT time-out. First time-out is detected when, at the beginning of a clock cycle, TSR[WIS] = 1. (This occurs on the second cycle after the 0→1 transition of the appropriate time-base bit while TSR[ENW] = 1 and TSR[WIS] = 0.) The PPC405 immediately takes the interrupt. The address of the next sequential instruction is saved in SRR2; simultaneously, the contents of the MSR are written into SRR3 and the MSR is written with the values shown in Table 5-19. The high-order 16 bits of the program counter are then written with the contents of the EVPR and the low-order 16 bits of the program counter are written with 0x1020. Interrupt processing begins at the address in the program counter.

To clear the WDT interrupt, the interrupt handling routine must clear the WDT interrupt bit TSR[WIS]. Clearing is done by writing a word to TSR (using **mtspr**), with a 1 in any bit position that is to be cleared and 0 in all other bit positions. The data written to the status register is not direct data, but a mask; a 1 causes the bit to be cleared, and a 0 has no effect.

Executing the return from critical interrupt instruction (**rftci**) restores the contents of the program counter and the MSR from SRR2 and SRR3, respectively, and the PPC405 resumes execution at the contents of the program counter.

Table 5-19. Register Settings during Watchdog Timer Interrupts

SRR2	Written with the address of the next sequential instruction
SRR3	Written with the contents of the MSR
MSR	AP, APE, WE, CE, EE, PR, FP, FE0, DE, DWE, FE1, IR, DR ← 0 ME ← unchanged
PC	EVPR[0:15] 0x1020
TSR	WIS ← 1

5.19 Data TLB Miss Interrupt

The data TLB miss interrupt is generated if data translation is enabled and a valid TLB entry matching the EA and PID is not present. The address of the instruction generating the untranslatable effective data address is saved in SRR0. In addition, the hardware also saves the data address (that missed in the TLB) in the DEAR.

The ESR is set to indicate whether the excepting operation was a store (includes **dcbz**, **dcbi**, **dccci**).

The interrupt is precise. Program flow vectors to EVPR[0:15] || 0x1100.

The following registers are modified to the values specified in Table 5-20.

Table 5-20. Register Settings during Data TLB Miss Interrupts

SRR0	Set to the address of the instruction generating the effective address for which no valid translation exists.
SRR1	Set to the value of the MSR at the time of the interrupt
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x1100
DEAR	Set to the effective address of the failed access
ESR	DST ← 1 if excepting operation is a store operation (includes dcbi , dcbz , and dccci). MCI ← unchanged All other bits are cleared.

Programming Note: Data TLB miss interrupts can happen whenever data translation is enabled. Therefore, ensure that SRR0 and SRR1 are saved before enabling translation in an interrupt handler.

5.20 Instruction TLB Miss Interrupt

The instruction TLB miss interrupt is generated if instruction translation is enabled and execution is attempted for an instruction for which a valid TLB entry matching the EA and PID for the instruction fetch is not present. The instruction whose fetch caused the TLB miss is saved in SRR0.

The interrupt is precise with respect to the attempted execution of the instruction. Program flow vectors to EVPR[0:15] || 0x1200.

The following are modified to the values specified in Table 5-21.

Table 5-21. Register Settings during Instruction TLB Miss Interrupts

SRR0	Set to the address of the instruction for which no valid translation exists.
SRR1	Set to the value of the MSR at the time of the interrupt
MSR	AP, APE, WE, EE, PR, FP, FE0, DWE, FE1, IR, DR ← 0 CE, ME, DE ← unchanged
PC	EVPR[0:15] 0x1200

Programming Note: Instruction TLB miss interrupts can happen whenever instruction translation

is active. Therefore, insure that SRR0 and SRR1 are saved before enabling translation in an interrupt handler.

5.21 Debug Interrupt

Debug interrupts can be either *synchronous* or *asynchronous*. These debug events generate synchronous interrupts: branch taken (BT), data address compare (DAC), data value compare (DVC), instruction address compare (IAC), instruction completion (IC), and trap instruction (TIE). The exception (EXC) and unconditional (UDE) debug events generate asynchronous interrupts. See “Debug Events” on page 8-10 for more information about debug events.

For debug events, SRR2 is written with an address, which varies with the type of debug event, as shown in Table 5-22.

Table 5-22. SRR2 during Debug Interrupts

Debug Event	Address Saved in SRR2
BT DAC IAC TIE	Address of the instruction causing the event
DVC IC	Address of the instruction <i>following</i> the instruction that causing the event
EXC	Interrupt vector address of the initial exception that caused the exception debug event
UDE	Address of next instruction to be executed at time of UDE

SRR3 is written with the contents of the MSR and the MSR is written with the values shown in Table 5-23 on page 5-26. The high-order 16 bits of the program counter are then written with the contents of the EVPR; the low-order 16 bits of the program counter are written with 0x2000. Interrupt processing begins at the address in the program counter.

Executing an **rfci** instruction restores the program counter from SRR2 and the MSR from SRR3, and execution resumes at the address in the program counter.

Table 5-23. Register Settings during Debug Interrupts

SRR2	Written with an address as described in Table 5-22
SRR3	Written with the contents of the MSR
MSR	AP, APE, WE, CE, EE, PR, FP, FE0, DE, DWE, FE1, IR, DR ← 0 ME ← unchanged
PC	EVPR[0:15] 0x2000
DBSR	Set to indicate type of debug event.

Chapter 6. Timer Facilities

The PPC405 provides four timer facilities: a time base, a Programmable Interval Timer (PIT), a fixed interval timer (FIT), and a watchdog timer. The PIT is a Special Purpose Register (SPR). These facilities, which are driven by the same base clock, can, among other things, be used for:

- Time-of-day functions
- Data logging functions
- Peripherals requiring periodic service
- Periodic task switching

Additionally, the watchdog timer can help a system to recover from faulty hardware or software.

Figure 6-1 shows the relationship of the timers and the clock source to the time base.

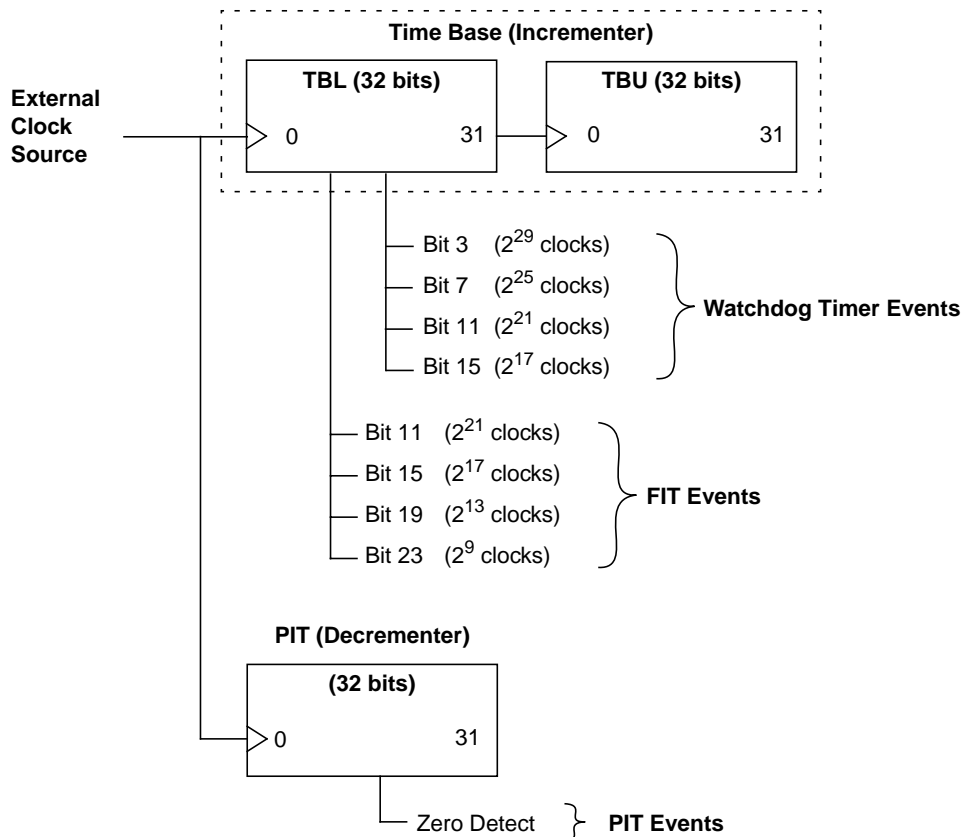


Figure 6-1. Relationship of Timer Facilities to the Time Base

6.1 Time Base

The PPC405 implements a 64-bit time base as required in *The PowerPC Architecture*. The time base, which increments once during each period of the source clock, provides a time reference. Read access to the time base is through the **mftb** instruction. **mftb** provides user-mode read-only access to

the time base. The TBR numbers (0x10C and 0x10D; TBL and TBU, respectively) that specify the time base registers to **mftb** are not SPR numbers. However, the PowerPC Architecture allows an implementation to handle **mftb** as **mfspr**. Accordingly, these register numbers cannot be used for other SPRs. PowerPC compilers cannot use **mftb** with register numbers other than those specified in the PowerPC Architecture as read-access time base registers (0x10C and 0x10D).

Write access to the time base, using **mtspr**, is privileged. Different register numbers are used for read access and write access. Writing the time base is accomplished by using SPR 0x11C and SPR 0x11D (TBL and TBU, respectively) as operands for **mtspr**.

The period of the 64-bit time base is approximately 2925 years for a 200 MHz clock source. The time base does not generate interrupts, even when it wraps. For most applications, the time base is set once at system reset and only read thereafter. Note that the FIT and the watchdog timer (discussed below) are driven by 0→1 transitions of bits from the TBL. Transitions caused by software alteration of TBL have the same effect as transitions caused by normal incrementing of the time base.

Figure 6-2 illustrates the TBL.



Figure 6-2. Time Base Lower (TBL)

0:31		Time Base Lower	Current count; low-order 32 bits of time base.
------	--	-----------------	--

Figure 6-3 illustrates the TBU.



Figure 6-3. Time Base Upper (TBU)

0:31		Time Base Upper	Current count; high-order 32 bits of time base.
------	--	-----------------	---

Table 6-1 summarizes the TBRs, instructions used to access the TBRs, and access restrictions.

Table 6-1. Time Base Access

Instructions		Register Number	Access Restrictions
TBU Upper 32 bits	mftbu RT <i>Extended mnemonic for mftb RT,TBU</i>	0x10D	Read-only
	mttbu RS <i>Extended mnemonic for mtspr TBU,RS</i>	0x11D	Privileged; write-only
TBL Lower 32 bits	mftb RT <i>Extended mnemonic for mftb RT,TBL</i>	0x10C	Read-only
	mttbl RS <i>Extended mnemonic for mtspr TBL,RS</i>	0x11C	Privileged; write-only

6.1.1 Reading the Time Base

The following code provides an example of reading the time base. **mftb** moves the low-order 32 bits of the time base to a GPR; **mftbu** moves the high-order 32 bits of the time base to a second GPR.

loop:

```

mftbu Rx          # load from TBU
mftb  Ry          # load from TBL
mftbu Rz          # load from TBU
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.

6.1.2 Writing the Time Base

The following code provides an example of writing the time base. Writing the time base is privileged. **mttbl** moves the contents of a GPR to the low-order 32 bits of the time base; **mttbu** moves the contents of a second GPR to the high-order 32 bits of the time base.

```

lwz   Rx, upper  # load 64-bit time base value into Rx and Ry
lwz   Ry, lower
li    Rz, 0
mttbl Rz          # force TBL to 0 to avoid rollover while writing TBU
mttbu Rx          # set TBU
mttbl Ry          # set TBL

```

6.2 Programmable Interval Timer (PIT)

The PIT is a 32-bit SPR that decrements at the same rate as the time base. The PIT is read and written using **mf spr** and **mt spr**, respectively. Writing to the PIT also simultaneously writes to a hidden reload register. Reading the PIT using **mf spr** returns the current PIT contents; the hidden reload register cannot be read. When a non-zero value is written to the PIT, it begins to decrement. A PIT event occurs when a decrement occurs on a PIT count of 1. When a PIT event occurs, the following occurs:

1. If the PIT is in auto-reload mode (the ARE field of the Timer Control Register (TCR) is 1), the PIT is loaded with the last value an **mt spr** wrote to the PIT. A decrement from a PIT count of 1 immediately causes a reload; no intermediate PIT content of 0 occurs.

If the PIT is not in auto-reload mode (TCR[ARE] = 0), a decrement from a PIT count of 1 simply causes a PIT content of 0.

2. TSR[PIS] is set to 1.
3. If enabled (TCR[PIE] = 1 and the EE field of the Machine State Register (MSR) is 1), a PIT interrupt is taken. See “Programmable Interval Timer (PIT) Interrupt” on page 5-22 for details of register behavior during a PIT interrupt.

The interrupt handler should use software to reset the PIS field of the Timer Status Register (TSR). This is done by using **mt spr** to write a word to the TSR having a 1 in TSR[PIS] and any other bits to be cleared, and a 0 in all other bits. The data written to the TSR is not direct data, but a mask. A 1 clears a bit; a 0 has no effect.

Using **mt spr** to force the PIT to 0 *does not* cause a PIT interrupt. However, decrementing that was ongoing at the instant of the **mt spr** instruction can cause the appearance of an interrupt. To eliminate the PIT as a source of interrupts, write a 0 to TCR[PIE], the PIT interrupt enable bit.

To eliminate all PIT activity:

1. Write a 0 to TCR[PIE]. This prevents PIT activity from causing interrupts.
2. Write a 0 to TCR[ARE]. This disables the PIT auto-reload feature.
3. Write zeroes to the PIT to halt PIT decrementing. Although this action does not cause a pit PIT interrupt to become pending, a near-simultaneous decrement to 0 might have done so.
4. Write a 1 to TSR[PIS] (PIT Interrupt Status bit). This clears TSR[PIS] to 0 (see “Timer Status Register (TSR)” on page 6-8). This also clears any pending PIT interrupt. Because the PIT stops decrementing, no further PIT events are possible.

If the auto-reload feature is disabled (TCR[ARE] = 0) when the PIT decrements to 0, the PIT remains 0 until software uses **mt spr** to reload it.

After a reset, TCR[ARE] = 0, which disables the auto-reload feature.

Figure 6-4 illustrates the PIT.

Figure 6-4. Programmable Interval Timer (PIT)

0:31		Programmed interval remaining	Number of clocks remaining until the PIT event
------	--	-------------------------------	--

6.2.1 Fixed Interval Timer (FIT)

The FIT provides timer interrupts having a repeatable period. The FIT is functionally similar to an auto-reload PIT, except that only a smaller fixed selection of interrupt periods are available.

The FIT exception occurs on 0→1 transitions of selected bits from the time base, as shown in Table 6-2.

Table 6-2. FIT Controls

TCR[FP]	TBL Bit	Period (Time Base Clocks)	Period (200 Mhz Clock)
0, 0	23	2^9 clocks	2.56 μ sec
0, 1	19	2^{13} clocks	40.96 μ sec
1, 0	15	2^{17} clocks	0.655 msec
1, 1	11	2^{21} clocks	10.49 msec

The TSR[FIS] field logs a FIT exception as a pending interrupt. A FIT interrupt occurs if TCR[FIE] and MSR[EE] are enabled at the time of the FIT exception. “Fixed Interval Timer (FIT) Interrupt” on page 5-23 describes register settings during a FIT interrupt.

The interrupt handler should reset TSR[FIS]. This is done by using **mtspr** to write a word to the TSR having a 1 in TSR[FIS] and any other bits to be cleared, and a 0 in all other bits. The data written to the TSR is not direct data, but a mask. A 1 clears a bit and a 0 has no effect.

6.3 Watchdog Timer

The watchdog timer aids system recovery from software or hardware faults.

A watchdog timeout occurs on 0→1 transitions of a selected bit from the time base, as shown in the following table.

Table 6-3. Watchdog Timer Controls

TCR[WP]	TBL Bit	Period (Time Base Clocks)	Period (200 MHz Clock)
0,0	15	2^{17} clocks	0.655 msec
0,1	11	2^{21} clocks	10.49 msec
1,0	7	2^{25} clocks	0.168 sec
1,1	3	2^{29} clocks	2.684 sec

If a watchdog timeout occurs while $TSR[WIS] = 0$ and $TSR[ENW] = 1$, a watchdog interrupt occurs if the interrupt is enabled by $TCR[WIE]$ and $MSR[CE]$. “Watchdog Timer” on page 6-6 describes register behavior during a watchdog interrupt.

The interrupt handler should reset the $TSR[WIS]$ bit. This is done by using **mtspr** to write a word to the TSR having a 1 in $TSR[WIS]$ and any other bits to be cleared, and a 0 in all other bits. The data written to the TSR is not direct data, but a mask. A 1 clears a bit and a 0 has no effect.

If a watchdog timeout occurs while $TSR[WIS] = 1$ and $TSR[ENW] = 1$, a hardware reset occurs if enabled by a non-zero value of $TCR[WRC]$. In other words, a reset can occur if a watchdog timeout occurs while a previous watchdog timeout is pending. The assumption is that $TSR[WIS]$ was not cleared because the processor could not execute the watchdog handler, leaving reset as the only way to restart the system. Note that after $TCR[WRC]$ is set to a non-zero value, it cannot be reset by software. This prevents errant software from disabling the watchdog timer reset capability. After a reset, the initial value of $TCR[WRC] = 00$.

Figure 6-5 describes the watchdog state machine. In the figure, numbers in parentheses refer to descriptions of operating modes that follow the table.

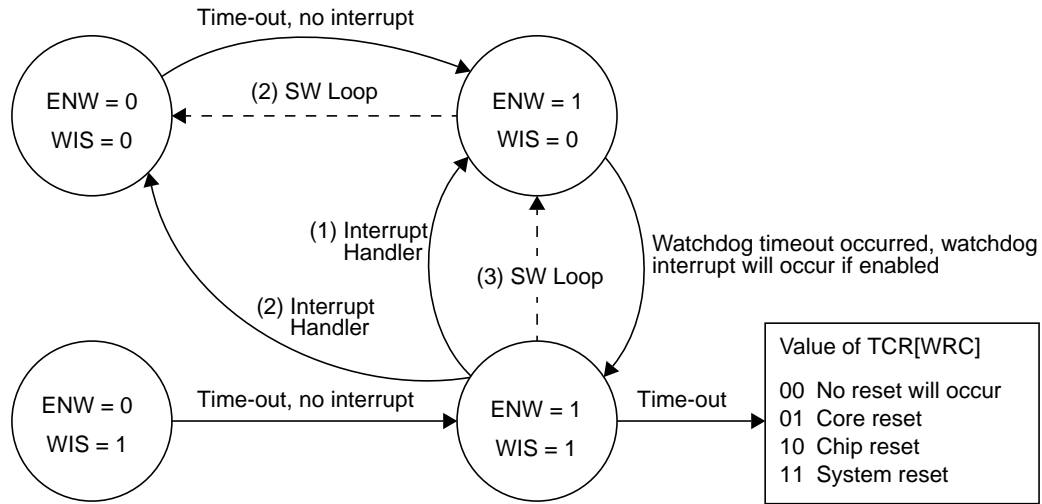


Figure 6-5. Watchdog Timer State Machine

Enable Next Watchdog TSR[ENW]	Watchdog Timer Status TSR[WIS]	Action When Timer Interval Expires
0	0	Set TSR[ENW] = 1.
0	1	Set TSR[ENW] = 1.
1	0	Set TSR[WIS] = 1. If TCR[WIE] = 1 and MSR[CE] = 1, then interrupt.
1	1	Cause the watchdog reset action specified by TCR[WRC]. On reset, copy current TCR[WRC] to TSR[WRS] and clear TCR[WRC], disabling the watchdog timer.

The controls described in Figure 6-5 imply three different ways of using the watchdog timer. The modes assume that TCR[WRC] was set to allow processor reset by the watchdog timer:

1. Always take a pending watchdog interrupt, and never attempt to prevent its occurrence. (This mode is described in the preceding text.)
 - a. Clear TSR[WIS] in the watchdog timer handler.
 - b. Never use TSR[ENW].
2. Always take a pending watchdog interrupt, but avoid it whenever possible by delaying a reset until a second watchdog timer occurs.

This assumes that a recurring code loop of known maximum duration exists outside the interrupt handlers, or that a FIT interrupt handler is operational. One of these mechanisms clears TSR[ENW] more frequently than the watchdog period.

- a. Clear TSR[ENW] to 0 in loop or in FIT interrupt handler.

To clear TSR[ENW], use **mtspr** to write a 1 to TSR[ENW] (and to any other bits that are to be cleared), with 0 in all other bit locations.

- b. Clear TSR[WIS] in watchdog timer handler.

It is not expected that a watchdog interrupt will occur every time, but only if an exceptionally high execution load delays clearing of TSR[ENW] in the usual time frame.

- 3. Never take a watchdog interrupt.

This assumes that a recurring code loop of reliable duration exists outside the interrupt handlers, or that a FIT interrupt handler is operational. This method only guarantees one watchdog timeout period before a reset occurs.

- a. Clear TSR[WIS] in the loop or in FIT handler.
- b. Never use TSR[ENW] but have it set.

6.4 Timer Status Register (TSR)

The TSR can be accessed for read or write-to-clear.

Status registers are generally set by hardware and read and cleared by software. The **mfspir** instruction reads the TSR. Clearing the TSR is performed by writing a word to the TSR, using **mtspr**, having a 1 in all fields to be cleared and a 0 in all other fields. The data written to the TSR is not direct data, but a mask. A 1 clears the field and a 0 has no effect.




Figure 6-6. Timer Status Register (TSR)

0	ENW	Enable Next Watchdog 0 Action on next watchdog event is to set TSR[ENW] = 1. 1 Action on next watchdog event is governed by TSR[WIS].	Software must reset TSR[ENW] = 0 after each watchdog timer event.
1	WIS	Watchdog Interrupt Status 0 No Watchdog interrupt is pending. 1 Watchdog interrupt is pending.	
2:3	WRS	Watchdog Reset Status 00 No Watchdog reset has occurred. 01 Core reset was forced by the watchdog. 10 Chip reset was forced by the watchdog. 11 System reset was forced by the watchdog.	
4	PIS	PIT Interrupt Status 0 No PIT interrupt is pending. 1 PIT interrupt is pending.	

5	FIS	FIT Interrupt Status 0 No FIT interrupt is pending. 1 FIT interrupt is pending.
6:31		Reserved

6.5 Timer Control Register (TCR)

The TCR controls PIT, FIT, and watchdog timer operation.

The TCR[WRC] field is cleared to 0 by all processor resets. (Chapter 3, “Initialization,” describes the types of processor reset.) This field is set only by software. However, hardware does not allow software to clear the field after it is set. After software writes a 1 to a bit in the field, that bit remains a 1 until any reset occurs. This prevents errant code from disabling the watchdog timer reset function.

All processor resets clear TCR[ARE] to 0, disabling the auto-reload feature of the PIT.




Figure 6-7. Timer Control Register (TCR)

0:1	WP	Watchdog Period 00 2^{17} clocks 01 2^{21} clocks 10 2^{25} clocks 11 2^{29} clocks
2:3	WRC	Watchdog Reset Control 00 No Watchdog reset will occur. 01 Core reset will be forced by the Watchdog. 10 Chip reset will be forced by the Watchdog. 11 System reset will be forced by the Watchdog. TCR[WRC] resets to 00. This field can be set by software, but cannot be cleared by software, except by a software-induced reset.
4	WIE	Watchdog Interrupt Enable 0 Disable watchdog interrupt. 1 Enable watchdog interrupt.
5	PIE	PIT Interrupt Enable 0 Disable PIT interrupt. 1 Enable PIT interrupt.
6:7	FP	FIT Period 00 2^9 clocks 01 2^{13} clocks 10 2^{17} clocks 11 2^{21} clocks

8	FIE	FIT Interrupt Enable 0 Disable FIT interrupt. 1 Enable FIT interrupt.	
9	ARE	Auto Reload Enable 0 Disable auto reload. 1 Enable auto reload.	Disables on reset.
10:31		Reserved	

Chapter 7. Memory Management

The PPC405 has a 4-gigabyte (GB) address space, which is presented as a flat address space. The PPC405 memory management unit (MMU) performs address translation and protection functions. With appropriate system software, the MMU supports:

- Translation of effective addresses to real addresses
- Independent enabling of instruction and data address translation and protection
- Page-level access control using the translation mechanism
- Software control of page replacement strategy
- Additional virtual-mode control of protection using zones
- Real-mode write protection

7.1 MMU Overview

The instruction and integer units generate 32-bit effective addresses (EAs) for instruction fetches and data accesses, respectively. Instruction EAs are generated for sequential instruction fetches, and for instruction fetches causing changes in program flow (branches and interrupts). Data EAs are generated for load/store and cache control instructions. The MMU translates EAs into real addresses; the instruction cache unit (ICU) and data cache unit (DCU) use real addresses to access memory.

The PPC405 MMU supports demand-paged virtual memory and other memory 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. Sufficient information is available to correct the fault and restart the faulting instruction.

The MMU divides storage into pages. A page represents the granularity of EA translation and protection controls. Eight page sizes (1KB, 4KB, 16KB, 64KB, 256KB, 1MB, 4MB, 16MB) are simultaneously supported. A valid entry for a page containing the EA to be translated must be in the translation lookaside buffer (TLB) for address translation to be performed. EAs for which no valid TLB entry exists cause TLB-miss interrupts.

7.2 Address Translation

Fields in the Machine State Register (MSR) control the use of the MMU for address translation. The instruction relocate (IR) field of the MSR controls translation for instruction accesses. The data relocate (DR) field of the MSR controls the translation mechanism for data accesses. These fields, specified independently, can be changed at any time by a program in supervisor state. Note that all interrupts clear MSR[IR, DR] and place the processor in the supervisor state. Subsequent discussion about translation and protection assumes that MSR[IR, DR] are set, enabling address translation.

The processor references memory when it fetches an instruction, and when it executes load/store, branch, and cache control instructions. Processor accesses to memory use EAs to reference a memory location. When translation is enabled, the EA is translated into a real address, as illustrated in Figure 7-1 on page 7-2. The ICU or DCU uses the real address for the access. (When translation is not enabled, the EA is already a real address.)

In address translation, the EA is combined with an 8-bit process ID (PID) to create a 40-bit virtual address. The virtual address is compared to all of the TLB entries. A matching entry supplies the real address for the storage reference. Figure 7-1 illustrates the process.

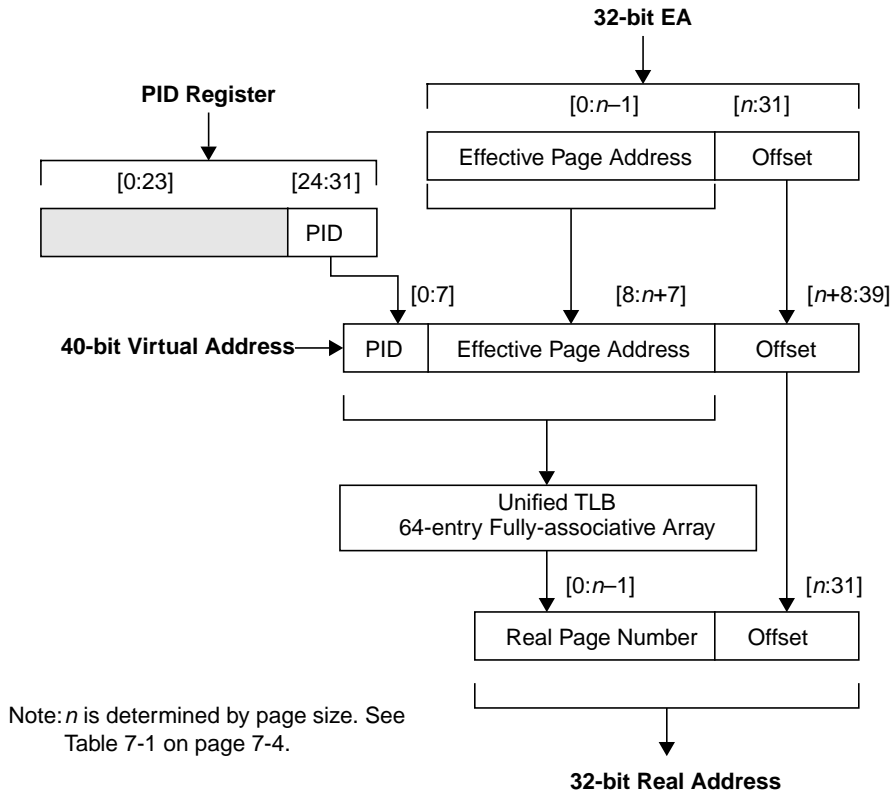


Figure 7-1. Effective to Real Address Translation Flow

7.3 Translation Lookaside Buffer (TLB)

The TLB is hardware that controls translation, protection, and storage attributes. The instruction and data units share a unified fully-associative TLB, in which any page entry (TLB entry) can be placed anywhere in the TLB. TLB entries are maintained under program control. System software determines the TLB entry replacement strategy and the format and use of page state information. A TLB entry contains the information required to identify the page, to specify translation and protection controls, and to specify the storage attributes.

7.3.1 Unified TLB

The unified TLB (UTLB) contains 64 entries; each has a TLBHI (tag) portion and a TLBLO (data) portion, as described in Figure 7-2 on page 7-3. TLBHI contains 36 bits; TLBLO contains 32 bits. When translation is enabled, the UTLB tag portion compares some or all of $EA_{0:21}$ with some or all of the effective page number $EPN_{0:21}$, based on the size bits $SIZE_{0:2}$. All 64 entries are simultaneously checked for a match. If an entry matches, the corresponding data portion of the UTLB provides the

real page number (RPN), access control bits (ZSEL, EX, WR), and storage attributes (W, I, M, G, E, U0




Figure 7-2. TLB Entries

The virtual address space is extended by adding an 8-bit translation ID (TID) loaded from the Process ID (PID) register during a TLB access. The PID identifies one of 255 unique software entities, usually used as a process or thread ID. TLBHI[TID] is compared to the PID during a TLB look-up.

Tag and data entries are written by copying data from GPRs and the PID, using the **tlbwe** instruction. Tag and data entries are read by copying data to GPRs and the PID, using the **tlbre** instruction. Software can search for specific entries using the **tlbsx** instruction.

7.3.2 TLB Fields

Each TLB entry describes a page that is enabled for translation and access controls. Fields in the TLB entry fall into four categories:

- Information required to identify the page to the hardware translation mechanism
- Control information specifying the translation
- Access control information
- Storage attribute control information

7.3.2.1 Page Identification Fields

When an EA is presented to the MMU for processing, the MMU applies several selection criteria to each TLB entry to select the appropriate entry. Although it is possible to place multiple entries into the TLB to match a specific EA and PID, this is considered a programming error, and the result of a TLB lookup for such an EA is undefined. The following fields in the TLB entry identify the page. Except as noted, all comparisons must succeed to validate an entry for subsequent use.

EPN (effective page number, 22 bits)

Compared to some number of the EA_{0:21} bits presented to the MMU. The number of bits corresponds to the page size.

The exact comparison depends on the page size, as shown in Table 7-1.

Table 7-1. TLB Fields Related to Page Size

Page Size	SIZE Field	<i>n</i> Bits Compared	EPN to EA Comparison	RPN Bits Set to 0
1KB	000	22	EPN _{0:21} ↔ EA _{0:21}	—
4KB	001	20	EPN _{0:19} ↔ EA _{0:19}	RPN _{20:21}
16KB	010	18	EPN _{0:17} ↔ EA _{0:17}	RPN _{18:21}
64KB	011	16	EPN _{0:15} ↔ EA _{0:15}	RPN _{16:21}
256KB	100	14	EPN _{0:13} ↔ EA _{0:13}	RPN _{14:21}
1MB	101	12	EPN _{0:11} ↔ EA _{0:11}	RPN _{12:21}
4MB	110	10	EPN _{0:9} ↔ EA _{0:9}	RPN _{10:21}
16MB	111	8	EPN _{0:7} ↔ EA _{0:7}	RPN _{8:21}

SIZE (page size, 3 bits)

Selects one of the eight page sizes, 1KB–16MB, listed in Table 7-1.

V (valid, 1 bit)

Indicates whether a TLB entry is valid and can be used for translation.

A valid TLB entry implies read access, unless overridden by zone protection. TLB_entry[V] can be written using a **tlbwe** instruction. The **tlbia** instruction invalidates all TLB entries.

TID (translation ID, 8 bits)

Loaded from the PID register during a **tlbwe** operation. The TID value is compared with the PID value during a TLB access. The TID provides a convenient way to associate a translation with one of 255 unique software entities, typically a process or thread ID maintained by operating system software. Setting TLBHI_entry[TID] = 0x00 disables TID-PID comparison and identifies a TLB entry as valid for all processes; the value of the PID register is then irrelevant.

7.3.2.2 Translation Field

When a TLB entry is identified as matching an EA (and possibly the PID), TLBLO_entry[RPN] defines how the EA is translated.

RPN (real page number, 22 bits)

Replaces some, or all, of EA_{0:21}, depending on page size. For example, a 16KB page uses EA_{0:17} for comparison. The translation mechanism replaces EA_{0:17} with TLBLO_entry[RPN]_{0:17} to form the physical address, and EA_{18:31} becomes the real page offset, as illustrated in Figure 7-1.

Programming Note: Software must set all unused bits of RPN (as determined by page size) to 0. See Table 7-1.

7.3.2.3 Access Control Fields

Several access controls are available in the UTLB entries.

ZSEL (zone select, 4 bits)

Selects one of 16 zone fields (Z0—Z15) from the Zone Protection Register (ZPR). The ZPR field bits can modify the access protection specified by the TLB_entry[V, EX, WR] bits of a TLB entry. Zone protection is described in detail in “Zone Protection” on page 7-14.

EX (execute enable, 1 bit)

When set (TLBLO_entry[EX] = 1), enables instruction execution at addresses within a page. ZPR settings can override TLBLO_entry[EX]; see “Zone Protection” on page 7-14, for more information.

WR (write-enable 1 bit)

When set (TLBLO_entry[WR] = 1), enables store operations to addresses in a page. ZPR settings can override TLBLO_entry[WR]; see “Zone Protection” on page 7-14.

7.3.2.4 Storage Attribute Fields

TLB entries contain bits that control and provide information about the storage control attributes. Four of the attributes (W, I, M, and G) are defined in the PowerPC Architecture. The E storage attribute is defined in the IBM PowerPC Embedded Environment. The U0 attribute is implementation-specific.

W (write-through, 1 bit)

When set (TLBLO_entry[W] = 1), stores are specified as write-through. If data in the referenced page is in the data cache, a store updates the cached copy of the data and the external memory location. Contrast this with a write-back strategy, which updates memory only when a cache line is flushed.

In real mode, the Data Cache Write-through Register (DCWR) controls the write strategy.

Note that the PowerPC Architecture does not support memory models in which write-through is enabled and caching is inhibited. It is considered a programming error to use these memory models; the results are undefined.

I (caching inhibited, 1 bit)

When set (TLBLO_entry[I] = 1), a memory access is completed by using the location in main memory, bypassing the cache arrays. During the access, the accessed location is not put into the cache arrays.

In real mode, the Instruction Cache Cachability Register (ICCR) and Data Cache Cachability Register (DCCR) control cachability. In these registers, the setting of the bit is reversed; 1 indicates that a storage control region is cachable, rather than caching inhibited.

Note that the PowerPC Architecture does not support memory models in which write-through is enabled and caching is inhibited. It is considered a programming error to use these memory models; the results are undefined.

It is considered a programming error if the target location of a load/store, **dcbz**, or fetch access to caching inhibited storage is in the cache; the results are undefined. It is *not* considered a programming error for the target locations of other cache control instructions to be in the cache when caching is inhibited.

M (memory coherent, 1 bit)

For implementations that support multiprocessing, the M storage attribute improves the performance of memory coherency management. Because the PPC405 does not provide multi-processor support or hardware support for data coherency, the M bit is implemented, but has no effect.

G (guarded, 1 bit)

When set ($TLBLO_entry[G] = 1$), indicates that the hardware cannot speculatively access the location for pre-fetching or out-of-order load access. The G storage attribute is typically used to protect memory-mapped I/O from inadvertent access. Attempted execution of an instruction from a guarded data storage address while instruction address translation is enabled results in an instruction storage interrupt because data storage and memory mapped I/O (MMIO) addresses are not used to contain instructions.

An instruction fetch from a guarded region does not occur until the execution pipeline is empty, thus guaranteeing that the access is necessary and therefore not speculative. For this reason, performance is degraded when executing out of guarded regions, and software should avoid unnecessarily marking regions of instruction storage as guarded.

In real mode, the Storage Guarded Register (SGR) controls guarding.

U0 (user-defined attribute, 1 bit)

When set ($TLBLO[U0] = 1$), indicates the user-defined attribute applies to the data in the associated page.

In real mode, the Storage User-defined 0 Register (SU0R) controls the setting of the U0 storage attribute.

E (endian, 1 bit)

When set ($TLBLO[E] = 1$), indicates that data in the associated page is stored in true little endian format.

In real mode, the Storage Little-Endian Register (SLER) controls the setting of the E storage attribute.

7.3.3 Shadow Instruction TLB

To enhance performance, four instruction-side TLB entries are kept in a four-entry fully-associative shadow array. This array, called the instruction TLB (ITLB), helps to avoid TLB contention between instruction accesses to the TLB and load/store operations. Replacement and invalidation of the ITLB entries is managed by hardware. See “Shadow TLB Consistency” on page 7-7 for details.

The ITLB can be considered a level-1 instruction-side TLB; the UTLB serves as the level-2 instruction-side TLB. The ITLB is used only during instruction fetches for storing instruction address translations. Each ITLB entry contains the translation information for a page. The processor uses the ITLB for address translation of instruction accesses when $MSR[IR] = 1$.

7.3.3.1 ITLB Accesses

The instruction unit accesses the ITLB independently of the rest of the MMU. ITLB accesses are transparent to the executing program, except that ITLB hits contribute to higher overall instruction throughput by allowing data address translations to occur in parallel. Therefore, when instruction accesses hit in the ITLB, the address translation mechanisms in the UTLB are available for use by data accesses simultaneously.

The ITLB requests a new entry from the UTLB when an ITLB miss occurs. A four-cycle latency occurs at each ITLB miss that is also a UTLB hit; the latency is longer if it is also a UTLB miss, or if there is contention for the UTLB from the data side. A round-robin replacement algorithm replaces existing entries with new entries.

7.3.4 Shadow Data TLB

To enhance performance, eight data-side TLB entries are kept in a eight-entry fully-associative shadow array. This array, called the data TLB (DTLB), helps to avoid TLB contention between instruction accesses to the TLB and load/store operations. Replacement and invalidation of the DTLB entries is managed by hardware. See “Shadow TLB Consistency” on page 7-7 for details.

The DTLB can be considered a level-1 data-side TLB; the UTLB serves as the level-2 data-side TLB. The DTLB is used only during instruction execute for storing data address translations. Each DTLB entry contains the translation information for a page. The processor uses the DTLB for address translation of data accesses when MSR[DR] = 1.

7.3.4.1 DTLB Accesses

The execute unit accesses the DTLB independently of the rest of the MMU. DTLB accesses are transparent to the executing program, except that DTLB hits contribute to higher overall instruction throughput by allowing instruction address translations to occur in parallel. Therefore, when data accesses hit in the DTLB, the address translation mechanisms in the UTLB are available for use by instruction accesses simultaneously.

The DTLB requests a new entry from the UTLB when a DTLB miss occurs. A three-cycle latency occurs at each DTLB miss that is also a UTLB hit; the latency is longer if it is also a UTLB miss. If there is contention for the UTLB from the instruction side, the data side has priority. A round-robin replacement algorithm replaces existing entries with new entries.

7.3.5 Shadow TLB Consistency

To help maintain the integrity of the shadow TLBs, the processor invalidates the ITLB and DTLB contents when the following context-synchronizing events occur:

- **isync** instruction
- Processor context switch (all interrupts, **rfi**, **rfdi**)
- **sc** instruction

If software updates a translation/protection mechanism (UTLB, PID, ZPR, or MSR) and must synchronize these updates with the ITLB and DTLB, the *software* must perform the necessary context synchronization.

A typical example is the manipulation of the TLB by an operating system within an interrupt handler for a TLB miss. Upon entry to the interrupt handler, the contents of the ITLB and DTLB are invalidated

and translation is disabled. If the operating system simply made the TLB updates and returned from the handler (using **rfi** or **rftci**), no additional explicit software action would be required to synchronize the ITLB and DTLB.

If, instead, the operating system enables translation within the handler and then performs TLB updates within the handler, those updates would not be effective in the ITLB and DTLB until **rftci** is executed to return from the handler. For those TLB updates to be reflected in the ITLB and DTLB *within* the handler, an **isync** must be issued after TLB updates finish. Failure to properly synchronize the shadow TLBs can cause unexpected behavior.

Programming Note: As a rule of thumb, follow software manipulation of an translation mechanism (if performed while translation is active) with a context-synchronizing operation (usually **isync**).

Figure 7-3 illustrates the relationship of the shadow TLBs and UTLB in address translation:

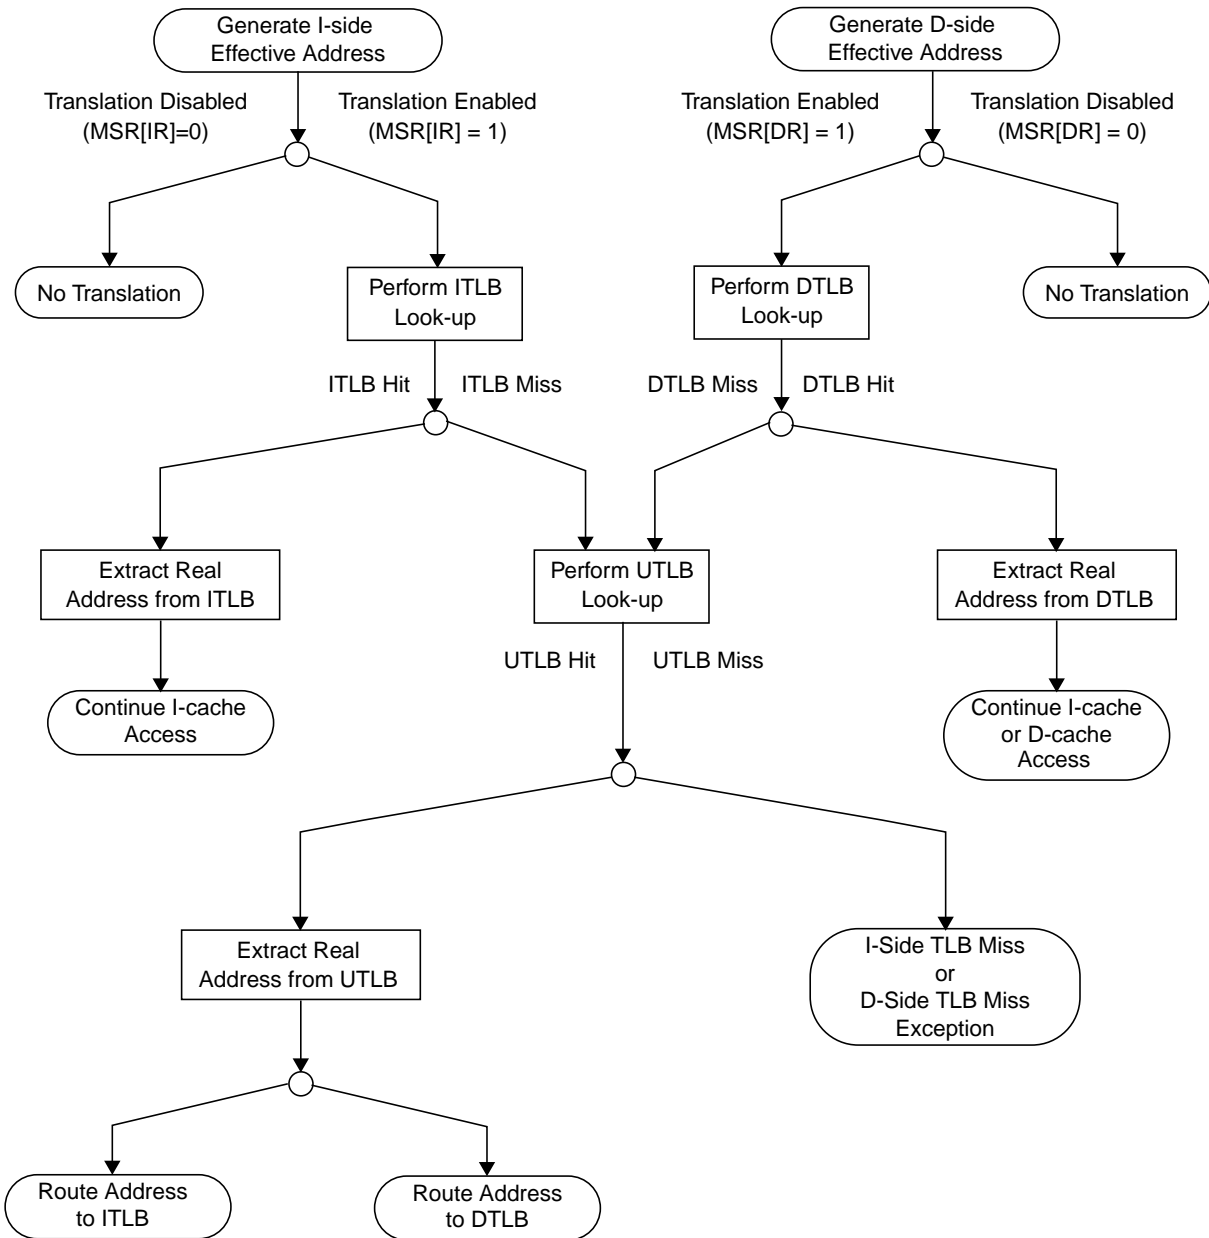


Figure 7-3. ITLB/DTLB/UTLB Address Resolution

7.4 TLB-Related Interrupts

The processor relies on interrupt handling software to implement paged virtual memory, and to enforce protection of specified memory pages.

When an interrupt occurs, the processor clears MSR[IR, DR]. Therefore, at the start of all interrupt handlers, the processor operates in real mode for instruction accesses and data accesses. Note that when address translation is disabled for an instruction fetch or load/store, the EA is equal to the real

address and is passed directly to the memory subsystem (including cache units). Such untranslated addresses bypass all memory protection checks that would otherwise be performed by the MMU.

When translation is enabled, MMU accesses can result in the following interrupts:

- Data storage interrupt
- Instruction storage interrupt
- Data TLB miss interrupt
- Instruction TLB miss interrupt
- Program interrupt

7.4.1 Data Storage Interrupt

A data storage interrupt is generated when data address translation is active, and the desired access to the EA is not permitted for one of the following reasons:

- In the problem state
 - **icbi**, load/store, **dcbz**, or **dcbf** with an EA whose zone field is set to no access ($ZPR[Zn] = 00$). In this case, **dcbt** and **dcbtst** no-op, rather than cause an interrupt. Privileged instructions cannot cause data storage interrupts.
 - Stores, or **dcbz**, to an EA having $TLB[WR] = 0$ (write access disabled) and $ZPR[Zn] \neq 11$. (The privileged instructions **dcbi** and **dccci** are treated as “stores”, but cause program interrupts, rather than data storage interrupts.)
- In supervisor state
 - Data store, **dcbi**, **dcbz**, or **dccci** to an EA having $TLB[WR] = 0$ and $ZPR[Zn]$ other than 11 or 10.

dcba does not cause data storage exceptions (cache line locking or protection). If conditions occur that would otherwise cause such an exception, **dcba** is treated as a no-op.

“Zone Protection” on page 7-14 describes zone protection in detail. See “Data Storage Interrupt” on page 5-16 for a detailed discussion of the data storage interrupt.

7.4.2 Instruction Storage Interrupt

An instruction storage interrupt is generated when instruction address translation is active and the processor attempts to execute an instruction at an EA for which fetch access is not permitted, for any of the following reasons:

- In the problem state
 - Instruction fetch from an EA with $ZPR[Zn] = 00$.
 - Instruction fetch from an EA having $TLB_entry[EX] = 0$ and $ZPR[Zn] \neq 11$.
 - Instruction fetch from an EA having $TLB_entry[G] = 1$.
- In the supervisor state
 - Instruction fetch from an EA having $TLB_entry[EX] = 0$ and $ZPR[Zn]$ other than 11 or 10.
 - Instruction fetch from an EA having $TLB_entry[G] = 1$.

See “Zone Protection” on page 7-14 for a detailed discussion of zone protection. See “Instruction Storage Interrupt” on page 5-17 for a detailed discussion of the instruction storage interrupt.

7.4.3 Data TLB Miss Interrupt

A data TLB miss interrupt is generated if data address translation is enabled and a valid TLB entry matching the EA and PID is not present. The interrupt applies to data access instructions and cache operations (excluding cache touch instructions).

See “Data TLB Miss Interrupt” on page 5-25 for a detailed discussion.

7.4.4 Instruction TLB Miss Interrupt

The instruction TLB miss interrupt is generated if instruction address translation is enabled and execution is attempted for an instruction for which a valid TLB entry matching the EA and PID for the instruction fetch is not present.

See “Instruction TLB Miss Interrupt” on page 5-25 for a detailed discussion.

7.4.5 Program Interrupt

When the TIE_cpuMmuEn signal is tied to 0, the TLB instructions (**tlbia**, **tlbre**, **tlbsx**, **tlbsync**, and **tlbwe**) are treated as illegal instructions. When execution of any of these instructions occurs under this circumstance, a program interrupt results.

See “Program Interrupt” on page 5-20 for a detailed discussion.

When TIE_cpuMmuEn is tied to 0, MSR[IR,DR] = 0.

Programming Note: When TIE_cpuMmuEn is tied to 0, MSR[IR,DR] = 0 upon execution of an **rfi** or **rfci** instruction, even if an interrupt handler sets MSR[IR] = 1 or MSR[DR] = 1 in Save/Restore Register 0 (SRR0) or SRR3.

See “Program Interrupt” on page 5-20 for a detailed discussion.

7.5 TLB Management

The processor does not imply any format for the page tables or the page table entries because there is no hardware support for page table management. Software has complete flexibility in implementing a replacement strategy, because software does the replacing. For example, software can “lock” TLB entries that correspond to frequently used storage by electing to never replace them, so that those entries are never cast out of the TLB.

TLB management is performed by software with some hardware assist, consisting of:

- Storage of the missed EA in the Save/Restore Register 0 (SRR0) for an instruction-side miss, or in the Data Exception Address Register (DEAR) for a data-side miss.
- Instructions for reading, writing, searching, and invalidating the TLB, as described briefly in the following subsections. See Chapter 9, “Instruction Set,” for detailed instruction descriptions.

7.5.1 TLB Search Instructions (tlbsx/tlbsx.)

tlbsx locates entries in the TLB, to find the TLB entry associated with an interrupt, or to locate candidate entries to cast out. **tlbsx** searches the UTLB array for a matching entry. The EA is the value to be matched; $EA = (RA|0)+(RB)$.

If the TLB entry is found, its index is placed in $RT_{26:31}$. RT can then serve as the source register for a **tlbre** or **tlbwe** instruction to read or write the entry, respectively. If no match is found, the contents of RT are undefined.

tlbsx. sets the Condition Register (CR) bit $CR0_{EQ}$. The value of $CR0_{EQ}$ depends on whether an entry is found: $CR0_{EQ} = 1$ if an entry is found; $CR0_{EQ} = 0$ if no entry is found.

7.5.2 TLB Read/Write Instructions (tlbre/tlbwe)

TLB entries can be accessed for reading and writing by **tlbre** and **tlbwe**, respectively. Separate extended mnemonics are available for the TLBHI (tag) and TLBLO (data) portions of a TLB entry.

7.5.3 TLB Invalidate Instruction (tlbia)

tlbia sets $TLB_entry[V] = 0$ to invalidate all TLB entries. All other TLB entry fields remain unchanged.

Using **tlbwe** to set $TLB_entry[V] = 0$ invalidates a specific TLB entry.

7.5.4 TLB Sync Instruction (tlbsync)

tlbsync guarantees that all TLB operations have completed for all processors in a multi-processor system. PPC405 provides no multiprocessor support, so this instruction performs no function. The instruction is included to facilitate code portability.

7.6 Recording Page References and Changes

When system software manages virtual memory, the software views physical memory as a collection of pages. Each page is associated with at least one TLB entry. To manage memory effectively, system software often must know whether a particular page has been referenced or modified. Note that this involves more than knowing whether a particular TLB entry was used to reference or alter memory, because multiple TLB entries can translate to the same page.

When system software manages a demand-paged environment, and the software needs to replace the contents of a page with other data, previously referenced pages (accessed for any purpose) are more likely to be maintained than pages that were never referenced. If the contents of a page must be replaced, and data contained in that page was modified, system software generally must write the contents of the modified page to the backing store before replacing its contents. System software must maintain records to control the environment.

Similarly, when system software manages TLB entries, the software often must know whether a particular TLB entry was referenced. When the system software must select a TLB entry to cast out, previously referenced entries are more likely to be maintained than entries which were never referenced. System software must also maintain records for this purpose.

The PPC405 does not provide hardware reference or change bits, but TLB miss interrupts and data storage interrupts enable system software to maintain reference information for TLB entries and their associated pages, respectively.

A possible algorithm follows. First, the TLB entries are built, with each $TLB_entry[V, WR] = 0$. System software retains the index and EPN of each entry.

The first attempt by application code to access a page causes a TLB miss interrupt, because its TLB entry is marked invalid. The TLB miss handler records the reference to the TLB entry (and to the associated page) in a data structure, then sets $TLB_entry[V] = 1$. (Note that $TLB_entry[V]$ can be considered a reference bit for the TLB entry.) Subsequent read accesses to the page associated with the TLB entry proceed normally.

In the example just given for recording TLB entry references, the first write access to the page using the TLB entry, after the entry is made valid, causes a data storage interrupt because write access was turned off. The TLB miss handler records the write to the page in a data structure, for use as a “changed” flag, then sets $TLB_entry[WR] = 1$ to enable write access. (Note that $TLB_entry[WR]$ can be considered a change bit for the page.) Subsequent write accesses to the page proceed normally.

7.7 Access Protection

The PPC405 provides virtual-mode access protection. The TLB entry enables system software to control general access for programs in the problem state, and control write and execute permissions for all pages. The TLB entry can specify zone protection that can override the other access control mechanisms supported in the TLB entries.

TLB entry and zone protection methods also support access controls for cache operation and string loads/stores.

7.7.1 Access Protection Mechanisms in the TLB

For MMU access protection to be in effect, one or both of $MSR[IR]$ or $MSR[DR]$ must be set to one to enable address translation. $MSR[IR]$ enables protection on instruction fetches, which are inherently read-only. $MSR[DR]$ enables protection on data accesses (loads/stores).

7.7.1.1 General Access Protection

The translation ID ($TLB_entry[TID]$) provides the first level of MMU access protection. This 8-bit field, if non-zero, is compared to the contents of $TLB_entry[PID]$. These fields must match in a valid TLB entry if any access is to be allowed. In typical use, it is assumed that a program in the supervisor state, such as a real-time operating system, sets the PID before starting a problem state program that is subject to access control.

If TLB_entry[TID] = 0x00, the associated memory page is accessible to all programs, regardless of their PID. This enables multiple processes to share common code and data. The common area is still subject to all other access protection mechanisms. Figure 7-4 illustrates the PID.



Figure 7-4. Process ID (PID)

0:23		Reserved
24:31		Process ID

7.7.1.2 Execute Permissions

If instruction address translation is enabled, instruction fetches are subject to MMU translation and have MMU access protection. Fetches are inherently read-only, so write protection is not needed. Instead, using TLB_entry[EX], a memory page is marked as executable (contains instructions) or not executable (contains only data or memory-mapped control hardware).

If an instruction is pre-fetched from a memory page for which TLB_entry[EX] = 0, the instruction is tagged as an error. If the processor subsequently attempts to execute this instruction, an instruction storage interrupt results. This interrupt is precise with respect to the attempted execution. If the fetcher discards the instruction without attempting to execute it, no interrupt will result.

Zone protection can alter execution protection.

7.7.1.3 Write Permissions

If MSR[DR] = 1, data loads and stores are subject to MMU translation and are afforded MMU access protection. The existence of a TLB entry describing a memory page implies read access; write access is controlled by TLB_entry[WR].

If a store (including those caused by **dcbz**, **dcbi**, or **dccci**) is made to an EA having TLB_entry[WR] = 0, a data storage interrupt results. This interrupt is precise.

Zone protection can alter write protection (see “Zone Protection” on page 7-14). In addition, only zone protection can prevent read access of a page defined by a TLB entry.

7.7.1.4 Zone Protection

Each TLB entry contains a 4-bit zone select (ZSEL) field. A zone is an arbitrary identifier for grouping TLB entries (memory pages) for purposes of protection. As many as 16 different zones may be defined. Any zone can have any number of member pages.

Each zone is associated with a 2-bit field (Z0–Z15) in the ZPR. The values of the field define how protection is applied to all pages that are member of that zone. Changing the value of the ZPR field can alter the protection attributes of all pages in the zone. Without ZPR, the change would require finding, reading, altering, and rewriting the TLB entry for each page in a zone, individually. The ZPR provides a much faster means of altering the protection for groups of memory pages.

The ZSEL values 0–15 select ZPR fields Z0–Z15, respectively.

The fields are defined within the ZPR as follows:

While it is common for TLB_entry[EX, WR] to be identical for all member pages in a group, this is not required. The ZPR field alters the protection defined by TLB_entry[EX] and TLB_entry[WR], on a page-by-page basis, as shown in the ZPR illustration. An application program (presumed to be running in the problem state) can have execute and write permissions as defined by TLB_entry[EX] and TLB_entry[WR] for the individual pages, or no access (denies loads, as well as stores and execution), or complete access.




Figure 7-5. Zone Protection Register (ZPR)

0:1	Z0	TLB page access control for all pages in this zone. In the problem state (MSR[PR] = 1): 00 No access 01 Access controlled by applicable TLB_entry[EX, WR] 10 Access controlled by applicable TLB_entry[EX, WR] 11 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted In the supervisor state (MSR[PR] = 0): 00 Access controlled by applicable TLB_entry[EX, WR] 01 Access controlled by applicable TLB_entry[EX, WR] 10 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted 11 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted
2:3	Z1	See the description of Z0.
4:5	Z2	See the description of Z0.
6:7	Z3	See the description of Z0.
8:9	Z4	See the description of Z0.
10:11	Z5	See the description of Z0.
12:13	Z6	See the description of Z0.
14:15	Z7	See the description of Z0.
16:17	Z8	See the description of Z0.
18:19	Z9	See the description of Z0.
20:21	Z10	See the description of Z0.
22:23	Z11	See the description of Z0.
24:25	Z12	See the description of Z0.
26:27	Z13	See the description of Z0.
28:29	Z14	See the description of Z0.
30:31	Z15	See the description of Z0.

Setting $ZPR[Zn] = 00$ for a ZPR field is the only way to deny read access to a page defined by an otherwise valid TLB entry. $TLB_entry[EX]$ and $TLB_entry[WR]$ do not support read protection. Note that the **icbi** instruction is considered a load with respect to access protection; executed in user mode, it causes a data storage interrupt if $MSR[DR] = 1$ and $ZPR[Zn] = 00$ is associated with the EA.

For a given ZPR field value, a program in supervisor state always has equal or greater access than a program in the problem state. System software can never be denied read (load) access for a valid TLB entry.

7.7.2 Access Protection for Cache Control Instructions

Architecturally the instructions **dcba**, **dcbi**, and **dcbz** are treated as “stores” because they can change data, or cause loss of data by invalidating a dirty line (a modified cache block).

Table 7-2 summarizes the conditions under which the cache control instructions can cause data storage interrupts.

Table 7-2. Protection Applied to Cache Control Instructions

Instruction	Possible Data Storage interrupt	
	When $ZPR[Zn] = 00$	When $TLB_entry[WR] = 0$
dcba	No (instruction no-ops)	No (instruction no-ops)
dcbf	Yes	No
dcbi	No	Yes
dcbst	Yes	No
dcbt	No (instruction no-ops)	No
dcbtst	No (instruction no-ops)	No
dcbz	Yes	Yes
dccci	No	Yes
dcread	No	No
icbi	Yes	No
icbt	No (instruction no-ops)	No
iccci	No	No
icread	No	No

If data address translation is enabled, and write permission is denied ($TLB_entry[WR] = 0$), **dcbi** and **dcbz** can cause data storage interrupts. **dcbz** can cause a data storage interrupt when executed in the problem state and all access is denied ($ZPR[Zn] = 00$); **dcbi** cannot cause a data storage interrupt because it is a privileged instruction.

The **dcba** instruction enables “speculative” line establishment in the cache arrays; the established lines do not cause a line fill. Because the effects of **dcba** are speculative, interrupts that would otherwise result when $ZPR[Zn] = 00$ or $TLB_entry[WR] = 0$ do not occur. In such cases, **dcba** is treated as a no-op.

The **dccci** instruction can also be considered a “store” because it can change data by invalidating a dirty line; however, **dccci** is not address-specific (it affects an entire congruence class regardless of

the operand address of the instruction). To restrict possible damage from an instruction which can change data and yet avoids the protection mechanism, the **dccci** instruction is privileged.

If data address translation is enabled, **dccci** can cause data storage interrupts when $TLB_entry[WR] = 0$; the operand is treated as if it were address-specific. **dccci** cannot cause a data storage interrupt when $ZPR[Zn] = 00$, because it is a privileged instruction.

Because **dccci** can cause data storage and TLB -miss interrupts, use of **dccci** is not recommended when $MSR[DR] = 1$; if **dccci** is used. Note that the specific operand address can cause an interrupt.

Architecturally, **dcbt** and **dcbtst** are treated as “loads” because they do not change data; they cannot cause data storage interrupts when $TLB_entry[WR] = 0$.

The cache block touch instructions **dcbt** and **dcbtst** are considered “speculative” loads; therefore, if a data storage interrupt would otherwise result from the execution of **dcbt** or **dcbtst** when $ZPR[Zn] = 00$, the instruction is treated as a no-op and the interrupt does not occur. Similarly, TLB miss interrupts do not occur for these instructions.

Architecturally, **dcbf** and **dcbst** are treated as “loads”. Flushing or storing a line from the cache is not architecturally considered a “store” because a store was performed to update the cache, and **dcbf** or **dcbst** only update main memory. Therefore, neither **dcbf** nor **dcbst** can cause data storage interrupts when $TLB_entry[WR] = 0$. Because neither instruction is privileged, they can cause data storage interrupts when $ZPR[Zn] = 00$ and data address translation is enabled.

dcread is a “load” from a non-specific address, and is privileged. Therefore, it cannot cause data storage interrupts when $ZPR[Zn] = 00$ or $TLB_entry[WR] = 0$.

icbi and **icbt** are considered “loads” and cannot cause data storage interrupts when $TLB_entry[WR] = 0$. **icbi** can cause data storage interrupts when $ZPR[Zn] = 00$.

The **iccci** instruction cannot change data; an instruction cache line cannot be dirty. The **iccci** instruction is privileged and is considered a load. It does not cause data storage interrupts when $ZPR[Zn] = 00$ or $TLB_entry[WR] = 0$.

Because **iccci** can cause a TLB miss interrupt, using **iccci** is not recommended when data address translation is enabled; if it is used, note that the specific operand address can cause an interrupt.

icread is considered a “load” from a non-specific address, and is privileged. Therefore, it cannot cause data storage interrupts when $ZPR[Zn] = 00$ or $TLB_entry[WR] = 0$.

7.7.3 Access Protection for String Instructions

The **stswx** instruction with string length equal to 0 ($XER[TBC] = 0$) is a no-op.

When data address translation is enabled and the Transfer Byte Count (TBC) field of the Fixed Point Exception Register (XER) is 0, neither **lswx** nor **stswx** can cause TLB miss interrupts, or data storage interrupts when $ZPR[Zn] = 0$ or $TLB_entry[WR] = 0$.

7.8 Real-Mode Storage Attribute Control

The PowerPC Architecture and the PowerPC Embedded Environment define several SPRs to control the following storage attributes in real mode: W, I, G, U0, and E. Note that the U0 and E attributes are not defined in the PowerPC Architecture. The E attribute is defined in the IBM PowerPC Embedded Environment, and the U0 attribute is implementation-specific. No storage attribute control register is

implemented for the M storage attribute because the PPC405 does not provide multi-processor support or hardware support for data coherency.

These SPRs, called storage attribute control registers, control the various storage attributes when address translation is disabled. When address translation is enabled, these registers are ignored, and the storage attributes supplied by the TLB entry are used (see “TLB Fields” on page 7-3).

The storage attribute control registers divide the 4GB real address space into thirty-two 128MB regions. In a storage attribute control register, bit 0 controls the lowest addressed 128MB region, bit 1 the next higher-addressed 128MB region, and so on. $EA_{0:4}$ specify a storage control region.

For detailed information on the function of the storage attributes, see “Storage Attribute Fields” on page 7-5.

7.8.1 Storage Attribute Control Registers

Figure 7-6 shows a generic storage attribute control register. The storage attribute control registers have the same bit numbering and address ranges.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
---	---	---	---	---	---	---	---	---	---	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----	----

Figure 7-6. Generic Storage Attribute Control Register

Bit	Address Range	Bit	Address Range
0	0x0000 0000–0x07FF FFFF	16	0x8000 0000–0x87FF FFFF
1	0x0800 0000–0x0FFF FFFF	17	0x8800 0000–0x8FFF FFFF
2	0x1000 0000–0x17FF FFFF	18	0x9000 0000–0x97FF FFFF
3	0x1800 0000–0x1FFF FFFF	19	0x9800 0000–0x9FFF FFFF
4	0x2000 0000–0x27FF FFFF	20	0xA000 0000–0xA7FF FFFF
5	0x2800 0000–0x2FFF FFFF	21	0xA800 0000–0xAFFF FFFF
6	0x3000 0000–0x37FF FFFF	22	0xB000 0000–0xB7FF FFFF
7	0x3800 0000–0x3FFF FFFF	23	0xB800 0000–0xBFFF FFFF
8	0x4000 0000–0x47FF FFFF	24	0xC000 0000–0xC7FF FFFF
9	0x4800 0000–0x4FFF FFFF	25	0xC800 0000–0xCFFF FFFF
10	0x5000 0000–0x57FF FFFF	26	0xD000 0000–0xD7FF FFFF
11	0x5800 0000–0x5FFF FFFF	27	0xD800 0000–0xDFFF FFFF
12	0x6000 0000–0x67FF FFFF	28	0xE000 0000–0xE7FF FFFF
13	0x6800 0000–0x6FFF FFFF	29	0xE800 0000–0xEFFF FFFF
14	0x7000 0000–0x77FF FFFF	30	0xF000 0000–0xF7FF FFFF
15	0x7800 0000–0x7FFF FFFF	31	0xF800 0000–0xFFFF FFFF

7.8.1.1 Data Cache Write-through Register (DCWR)

The DCWR controls write-through policy (the W storage attribute) for the data cache unit (DCU). Write-through is not applicable to the instruction cache unit (ICU).

After any reset, all DCWR bits are set to 0, which establishes a write-back write strategy for all regions.

The PowerPC Architecture does not support memory models in which write-through is enabled and caching is inhibited.

7.8.1.2 Data Cache Cachability Register (DCCR)

The DCCR controls the I storage attribute for data accesses and cache management instructions. Note that the polarity of the bits in this register is opposite to that of the I attribute in the TLB; $DCCR[S_n] = 1$ enables caching, while $TLB_entry[I] = 1$ inhibits caching.

After any reset, all DCCR bits are set to 0. No memory regions are cachable. Before memory regions can be designated as cachable in the DCCR, it is necessary to execute the **dccci** instruction once for each congruence class in the DCU cache array. This procedure invalidates all congruence classes. The DCCR can then be reconfigured, and the DCU can begin normal operation.

The PowerPC Architecture does not support memory models in which write-through is enabled and caching is inhibited.

7.8.1.3 Instruction Cache Cachability Register (ICCR)

The ICCR controls the I storage attribute for instruction fetches. Note that the polarity of the bits in this register is opposite of that of the I attribute ($ICCR[S_n] = 1$ enables caching, while $TLB_entry[I] = 1$ inhibits caching).

After any reset, all ICCR bits are set to 0. No memory regions are cachable. Before memory regions can be designated as cachable in the ICCR, it is necessary to execute the **iccci** instruction. This procedure invalidates all congruence classes. The ICCR can then be reconfigured, and the ICU can begin normal operation.

7.8.1.4 Storage Guarded Register (SGR)

The SGR controls the G storage attribute for instruction and data accesses.

This attribute does not affect data accesses; the PPC405 does not perform speculative loads or stores.

After any reset, all SGR bits are set to 1, marking all storage as guarded. For best performance, system software should clear the guarded attribute of appropriate regions as soon as possible. If $MSR[IR] = 1$, the G attribute comes from the TLB entry. Attempting to execute from a guarded region in translate mode causes an instruction storage interrupt. See “Instruction Storage Interrupt” on page 5-17 for more information.

7.8.1.5 Storage User-defined 0 Register (SU0R)

The Storage User-defined 0 Register (SU0R) controls the user-defined (U0) storage attribute for instruction and data accesses.

After any reset, all SU0R bits are set to 0.

7.8.1.6 Storage Little-Endian Register (SLER)

The SLER controls the E storage attribute for instruction and data accesses.

This attribute determines the byte ordering of storage. “Byte Ordering” on page 2-17 provides a detailed description of byte ordering in the IBM PowerPC Embedded Environment.

After any reset, all SLER bits are set to 0 (big endian).

Chapter 8. Debugging

The debug facilities of the PPC405 include support for debug modes for debugging during hardware and software development, and debug events that allow developers to control the debug process. Debug registers control the debug modes and debug events. The debug registers are accessed through software running on the processor or through a JTAG debug port. The debug interface is the JTAG debug port. The JTAG debug port can also be used for board test.

The debug modes, events, controls, and interface provide a powerful combination of debug facilities for a wide range of hardware and software development tools.

8.1 Development Tool Support

The RISCWatch product from IBM is an example of a development tool that uses the external debug mode, debug events, and the JTAG debug port to implement a hardware and software development tool. The RISCTrace™ feature of RISCWatch is an example of a development tool that uses the real-time trace capability of the PPC405.

8.2 Debug Modes

The PPC405 supports the following debug modes, each of which supports a type of debug tool or debug task commonly used in embedded systems development:

- Internal debug mode, which supports ROM monitors
- External debug mode, which supports JTAG debuggers
- Debug wait mode, which supports processor stopping or stepping for JTAG debuggers while servicing interrupts
- Real-time trace mode, which supports trigger events for real-time tracing

Internal and external debug modes can be enabled simultaneously. Both modes are controlled by fields in Debug Control Register 0 (DBCR0). Real-time trace mode is available only if internal, external, and debug wait modes are disabled.

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 generate debug interrupts, 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 at a dedicated interrupt vector and an external communications path to debug software problems. This mode, used while the processor executes instructions, enables debugging of operating system or application programs.

In this mode, debugger software is accessed through a communications port, such as a serial port, external to the processor core.

To enable internal debug mode, the Debug Control Register 0 (DBCR0) field IDM is set to 1 (DBCR0[IDM] = 1). To enable debug interrupts, MSR[DE] = 1. A debug interrupt occurs on a debug event only if DBCR0[IDM] = 1 and MSR[DE] = 1.

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 cause the processor to become architecturally frozen. While the processor is frozen, normal instruction execution stops and architected processor resources can be accessed and altered. External bus activity continues in external debug mode.

The JTAG mechanism can pass instructions to the processor for execution, allowing a JTAG debugger to display and alter processor resources, including memory.

The JTAG mechanism prevents the occurrence of a privileged exception when a privileged instruction is executed while the processor is in user mode.

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.

Because external debug mode relies only on internal processor resources, it can be used to debug system hardware and software.

In this mode, access to the processor is through the JTAG debug port.

To enable external debug mode, DBCR0[EDM] = 1. To enable debug interrupts, MSR[DE] = 1. A debug interrupt occurs on a debug event only if DBCR0[EDM] = 1 and MSR[DE] = 1.

8.2.3 Debug Wait Mode

In debug wait mode, debug events cause the PPC405 to enter a state in which interrupts can be serviced while the processor appears to be stopped.

Debug wait mode provides access to architected processor resources in a manner similar to external debug mode, except that debug wait mode allows the servicing of interrupt handlers. It supports stopping, starting, and stepping the processor, setting hardware and software breakpoints, and monitoring processor status. In this mode, if a debug event caused the processor to become architecturally frozen, an interrupt causes the processor to run an interrupt handler and return to the architecturally frozen state upon returning from the interrupt handler. While the processor is frozen, normal instruction execution stops and architected processor resources can be accessed and altered. External bus activity continues in debug wait mode.

The processor enters debug wait mode when internal and external debug modes are disabled (DBCR0[IDM, EDM] = 0), debug wait mode is enabled (MSR[DWE] = 1), debug wait is enabled by the JTAG debugger, and a debug event occurs.

For example, while the PPC405 core is in debug wait mode, an external device might generate an interrupt that requires immediate service. The PPC405 core can service the interrupt (vector to an interrupt handler and execute the interrupt handler code) and return to the previous stopped state.

Debug wait mode relies only on internal processor resources, 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.

In this mode, access to the processor is through the JTAG debug port.

8.2.4 Real-time Trace Debug Mode

Real-time trace debug mode supports the generation of trigger events for tracing the instruction stream being executed out of the instruction cache in real-time. In this mode, debug events can be used to control the collection of trace information through the use of trigger event generation. The broadcast of trace information is independent of the use of debug events as trigger events. This mode does not alter the processor performance.

A trace event occurs when internal and external debug modes are disabled ($DBCR0[IDM, EDM] = 0$) and a debug event occurs.

When a trace event occurs, a trace device can capture trace signals that provide the instruction trace information. Most trace events generated from debug events are blocked when internal debug, external debug, or debug wait modes are enabled

8.3 Processor Control

The PPC405 provides the following debug functions for processor control. Not all facilities are available in all debug modes.

Instruction Step	The processor is stepped one instruction at a time, while stopped, using the JTAG debug port.
Instruction Stuff	While the processor is stopped, instructions can be stuffed into the processor and executed using the JTAG debug port.
Halt	The processor can be stopped by activating an external halt signal on an external event, such as a logic analyzer trigger. This signal freezes the processor architecturally. While frozen, normal instruction execution stops and architected processor resources can be accessed and altered using the JTAG debug port. Normal execution resumes when the halt signal is deactivated.
Stop	The processor can be stopped using the JTAG debug port. Activating a stop causes the processor to become architecturally frozen. While frozen, normal instruction execution stops and the architected processor resources can be accessed and altered using the JTAG debug port.
Reset	An external reset signal, the JTAG debug port, or DBCR0 can request core, chip, and system resets.
Debug Events	A debug event triggers a debug operation. The operation depends on the debug mode. For more information and a list of debug events, see “Debug Events” on page 8-10.
Freeze Timers	The JTAG debug port or DBCR0 can control timer resources. The timers can be enabled to run, freeze always, or freeze on a debug event.
Trap Instructions	The trap instructions tw and twi can be used, with debug events, to implement software breakpoints.

8.4 Processor Status

The processor execution status, exception status, and most recent reset can be monitored.

Execution Status The JTAG debug port can monitor processor execution status to determine whether the processor is stopped, waiting, or running.

Exception Status The JTAG debug port can monitor the status of pending synchronous exceptions.

Most Recent Reset The JTAG debug port or an **mf spr** instruction can be used to read the Debug Status Register (DBSR) to determine the type of the most recent reset.

8.5 Debug Registers

Several debug registers, available to debug tools running on the processor, are not intended for use by application code. Debug tools control debug resources such as debug events. Application code that uses debug resources can cause the debug tools to fail, as well as other unexpected results, such as program hangs and processor resets.

Application code should not use the debug resources, including the debug registers.

8.5.1 Debug Control Registers

The debug control registers (DBCR0 and DBCR1) can enable and configure debug events, reset the processor, control timer operation during debug events, enable debug interrupts, and set the processor debug mode.

8.5.1.1 Debug Control Register 0 (DBCR0)




Figure 8-1. Debug Control Register 0 (DBCR0)

0	EDM	External Debug Mode 0 Disabled 1 Enabled
1	IDM	Internal Debug Mode 0 Disabled 1 Enabled
2:3	RST	Reset 00 No action 01 Core reset 10 Chip reset 11 System reset Attention: Writing 01, 10, or 11 to this field causes a processor reset request.

4	IC	Instruction Completion Debug Event 0 Disabled 1 Enabled	
5	BT	Branch Taken Debug Event 0 Disabled 1 Enabled	
6	EDE	Exception Debug Event 0 Disabled 1 Enabled	
7	TDE	Trap Debug Event 0 Disabled 1 Enabled	
8	IA1	IAC 1 Debug Event 0 Disabled 1 Enabled	
9	IA2	IAC 2 Debug Event 0 Disabled 1 Enabled	
10	IA12	Instruction Address Range Compare 1–2 0 Disabled 1 Enabled	Registers IAC1 and IAC2 define an address range used for IAC address comparisons.
11	IA12X	Enable Instruction Address Exclusive Range Compare 1–2 0 Inclusive 1 Exclusive	Selects the range defined by IAC1 and IAC2 to be inclusive or exclusive.
12	IA3	IAC 3 Debug Event 0 Disabled 1 Enabled	
13	IA4	IAC 4 Debug Event 0 Disabled 1 Enabled	
14	IA34	Instruction Address Range Compare 3–4 0 Disabled 1 Enabled	Registers IAC3 and IAC4 define an address range used for IAC address comparisons.
15	IA34X	Instruction Address Exclusive Range Compare 3–4 0 Inclusive 1 Exclusive	Selects range defined by IAC3 and IAC4 to be inclusive or exclusive.
16	IA12T	Instruction Address Range Compare 1-2 Toggle 0 Disabled 1 Enable	Toggles range 12 inclusive, exclusive DBCR[IA12X] on debug event.
17	IA34T	Instruction Address Range Compare 3–4 Toggle 0 Disabled 1 Enable	Toggles range 34 inclusive, exclusive DBCR[IA34X] on debug event.
18:30		Reserved	

31	FT	Freeze timers on debug event 0 Timers not frozen 1 Timers frozen
----	----	--

8.5.1.2 Debug Control Register1 (DBCR1)




Figure 8-2. Debug Control Register 1 (DBCR1)

0	D1R	DAC1 Read Debug Event 0 Disabled 1 Enabled	
1	D2R	DAC 2 Read Debug Event 0 Disabled 1 Enabled	
2	D1W	DAC 1 Write Debug Event 0 Disabled 1 Enabled	
3	D2W	DAC 2 Write Debug Event 0 Disabled 1 Enabled	
4:5	D1S	DAC 1 Size 00 Compare all bits 01 Ignore lsb (least significant bit) 10 Ignore two lsbs 11 Ignore five lsbs	Address bits used in the compare: Byte address Halfword address Word address Cache line (8-word) address
6:7	D2S	DAC 2 Size 00 Compare all bits 01 Ignore lsb (least significant bit) 10 Ignore two lsbs 11 Ignore five lsbs	Address bits used in the compare: Byte address Halfword address Word address Cache line (8-word) address
8	DA12	Enable Data Address Range Compare 1:2 0 Disabled 1 Enabled	Registers DAC1 and DAC2 define an address range used for DAC address comparisons
9	DA12X	Data Address Exclusive Range Compare 1:2 0 Inclusive 1 Exclusive	Selects range defined by DAC1 and DAC2 to be inclusive or exclusive
10:11		Reserved	

12:13	DV1M	Data Value Compare 1 Mode 00 Undefined 01 AND 10 OR 11 AND-OR	Type of data comparison used: All bytes selected by DBCR1[DV1BE] must compare to the appropriate bytes of DVC1. One of the bytes selected by DBCR1[DV1BE] must compare to the appropriate bytes of DVC1. The upper halfword or lower halfword must compare to the appropriate halfword in DVC1. When performing halfword compares set DBCR1[DV1BE] = 0011, 1100, or 1111.
14:15	DV2M	Data Value Compare 2 Mode 00 Undefined 01 AND 10 OR 11 AND-OR	Type of data comparison used All bytes selected by DBCR1[DV2BE] must compare to the appropriate bytes of DVC2. One of the bytes selected by DBCR1[DV2BE] must compare to the appropriate bytes of DVC2. The upper halfword or lower halfword must compare to the appropriate halfword in DVC2. When performing halfword compares set DBCR1[DV2BE] = 0011, 1100, or 1111.
16:19	DV1BE	Data Value Compare 1 Byte 0 Disabled 1 Enabled	Selects which data bytes to use in data value comparison
20:23	DV2BE	Data Value Compare 2 Byte 0 Disabled 1 Enabled	Selects which data bytes to use in data value comparison
24:31		Reserved	

8.5.2 Debug Status Register (DBSR)

The DBSR contains status on debug events and the most recent reset; the status is obtained by reading the DBSR. The status bits are normally set by debug events or by any of the three reset types.

Clearing DBSR fields is performed by writing a word to the DBSR, using the **mtdbsr** extended mnemonic, having a 1 in all bit positions to be cleared and a 0 in the all other bit positions. The data written to the DBSR is not direct data, but a mask. A 1 clears the bit and a 0 has no effect.

Application code should not use the DBSR.




Figure 8-3. Debug Status Register (DBSR)

0	IC	Instruction Completion Debug Event 0 Event did not occur 1 Event occurred
1	BT	Branch Taken Debug Event 0 Event did not occur 1 Event occurred
2	EDE	Exception Debug Event 0 Event did not occur 1 Event occurred
3	TIE	Trap Instruction Debug Event 0 Event did not occur 1 Event occurred
4	UDE	Unconditional Debug Event 0 Event did not occur 1 Event occurred
5	IA1	IAC1 Debug Event 0 Event did not occur 1 Event occurred
6	IA2	IAC2 Debug Event 0 Event did not occur 1 Event occurred
7	DR1	DAC1 Read Debug Event 0 Event did not occur 1 Event occurred
8	DW1	DAC1 Write Debug Event 0 Event did not occur 1 Event occurred
9	DR2	DAC2 Read Debug Event 0 Event did not occur 1 Event occurred
10	DW2	DAC2 Write Debug Event 0 Event did not occur 1 Event occurred
11	IDE	Imprecise Debug Event 0 No circumstance that would cause a debug event (if MSR[DE] = 1) occurred 1 A debug event would have occurred, but debug exceptions were disabled (MSR[DE] = 0)

12	IA3	IAC3 Debug Event 0 Event did not occur 1 Event occurred	
13	IA4	IAC4 Debug Event 0 Event did not occur 1 Event occurred	
14:21		Reserved	
22:23	MRR	Most Recent Reset 00 No reset has occurred since last cleared by software. 01 Core reset 10 Chip reset 11 System reset	This field is set to a value, indicating the type of reset, when a reset occurs.
24:31		Reserved	

8.5.3 Instruction Address Compare Registers (IAC1–IAC4)

The PPC405 can take a debug event upon an attempt to execute an instruction from an address. The address, which must be word-aligned, is defined in an IAC register. The DBCR0[IA1, IA2] fields of DBCR0 controls the instruction address compare (IAC) debug event.



Figure 8-4. Instruction Address Compare Registers (IAC1–IAC4)

0:29		Instruction Address Compare word address	Omit two low-order bits of complete address.
30:31		Reserved	

8.5.4 Data Address Compare Registers (DAC1–DAC2)

The PPC405 can take a debug event upon storage or cache references to addresses specified in the DAC registers. The specified addresses in the DAC registers are EAs of operands of storage references or cache instructions. The fields DBCR1[D1R], [D2R] and DBCR[D1W], [D2W] control the DAC-read and DAC-write debug events, respectively.

Addresses in the DAC registers specify exact byte EAs for DAC debug events. However, one may want to take a debug event on any byte within a halfword (ignore the least significant bit (LSb) of the DAC), on any byte within a word (ignore the two LSbs of DAC), or on any byte within eight words (ignore four LSbs of DAC). DBCR1[D1S, D2S] control the addressing options.

Errors related to execution of storage reference or cache instructions prevent DAC debug events.

0	31
---	----

Figure 8-5. Data Address Compare Registers (DAC1–DAC2)

0:31		Data Address Compare (DAC) byte address	DBCR0[D1S] determines which address bits are examined.
------	--	---	--

8.5.5 Data Value Compare Registers (DVC1–DVC2)

The PPC405 can take a debug event upon storage or cache references to addresses specified in the DAC registers, that also require the data at that address to match the value specified in the DVC registers. The data address compare for a DVC events works the same as for a DAC event. Cache operations do not cause DVC events. If the data at the address specified matches the value in the corresponding DVC register a DVC event will occur. The fields DBCR1[DV1M, DV2M] control how the data value are compared.

Errors related to execution of storage reference or cache instructions prevent DVC debug events.

0	31
---	----

Figure 8-6. Data Value Compare Registers (DVC1–DVC2)

0:31		Data Value to Compare
------	--	-----------------------

8.5.6 Debug Events

Debug events, enabled and configured by DBCR0 and DBCR1 and recorded in the DBSR, cause debug operations. A debug event occurs when an event listed in Table 8-1 on page 8-11 is detected. The debug operation is performed after the debug event.

In internal debug mode, the processor generates a debug interrupt when a debug event occurs. In external debug mode, the processor stops when a debug event occurs. When internal and external debug mode are both enabled, the processor stops on a debug event with the debug interrupt pending. When external and internal debug mode are both disabled, and debug wait mode is enabled the processor stops, but can be restarted by an interrupt. When all debug modes are disabled, debug events are recorded in the DBSR, but no action is taken.

Table 8-1 lists the debug events and the related fields in DBCR0, DBCR1, and DBSR. DBCR0 and DBCR1 enable the debugs events, and the DBSR fields report their occurrence.

Table 8-1. Debug Events

Event	Enabling DBCR0, DBCR1 Fields	Reporting DBSR Fields	Description
Instruction Completion	IC	IC	Occurs after completion of an instruction.
Branch Taken	BT	BT	Occurs before execution of a branch instruction determined to be taken.
Exception Taken	EDE	EXC	Occurs after an exception.
Trap Instruction	TDE	TIE	Occurs before execution of a trap instruction where the conditions are such that the trap will occur.
Unconditional	UDE	UDE	Occurs immediately upon being set by the JTAG debug port or the XXX_cpuUncondDebugEvent signal.
Instruction Address Compare	IA1, IA2, IA3, IA4, IA12, IA12X, IA12T, IA34, IA34X, IA34T	IA1, IA2, IA3, IA4	Occurs before execution of an instruction at an address that matches an address defined by the Instruction Address Compare Registers (IAC1–IAC4).
Data Address Compare	D1R, D1W, D1S, D2R, D2W, D2S, DA12, DA12X	DR2, DW2	Occurs before execution of an instruction that accesses a data address that matches the contents of the specified DAC register.
Data Value Compare	DV1M, DV2M, DV1BE, DV2BE	DR1, DW1	Occurs after execution of an instruction that accesses a data address for which a DAC occurs, and for which the value at the address matches the value in the specified DVC register.
Imprecise		IDE	Indicates that another debug event occurred while MSR[DE] = 0

8.5.7 Instruction Complete Debug Event

This debug event occurs after the completion of an instruction. If DBCR0[IDM] = 1, DBCR0[EDM] = 0 and MSR[DE] = 0 this debug event is disabled.

8.5.8 Branch Taken Debug Event

This debug event occurs before execution of a branch instruction determined to be taken. If DBCR0[IDM] = 1, DBCR0[EDM] = 0 and MSR[DE] = 0 this debug event is disabled.

8.5.9 Exception Taken Debug Event

This debug event occurs after an exception. Exception debug events always include the non-critical class of exceptions. When DBCR0[IDM] = 1 and DBCR0[EDM] = 0 the critical exceptions are not included.

8.5.10 Trap Taken Debug Event

This debug event occurs before execution of a trap instruction where the conditions are such that the trap will occur. When trap is enabled for a debug event, external debug mode is enabled, internal debug mode is enabled with MSR[DE] enabled, or debug wait mode is enabled, a trap instruction will not cause a program exception.

8.5.11 Unconditional Debug Event

This debug event occurs immediately upon being set by the JTAG debug port or the XXX_cpuUncondDebugEvent signal.

8.5.12 IAC Debug Event

This debug event occurs before execution of an instruction at an address that matches an address defined by the Instruction Address Compare Registers (IAC1–IAC4). DBCR0[IA1, IA2, IA3, IA4] enable IAC debug events IAC can be defined as an exact address comparison to one of the IAC n registers or on a range of addresses to compare defined by a pair of IAC n registers.

8.5.12.1 IAC Exact Address Compare

In this mode each IAC n register specifies an exact address to compare. These are enabled by setting DBCR0[IA n] = 1 and disabling IAC range compare (DBCR0[IA12X] = 0 for IAC1 and IAC2 and DBCR0[IA23X] = 0 for IAC3 and IAC4). The corresponding DBSR[IA n] bit displays the results of the debug event.

8.5.12.2 IAC Range Address Compare

In this mode a pair of IAC n registers are used to define a range of addresses to compare:

- Range 1:2 corresponds to IAC1 and IAC2
- Range 3:4 corresponds to IAC3 and IAC4

To enable Range 1:2, DBCR0[IA12] = 1 and DBCR0[IA1] or DBCR0[IA2] =1. An IAC event will be seen on the DBSR[IA n] field that corresponds to the enabled DBCR0[IA n] field. If DBCR0[IA1] and DBCR0[IA2] are enabled, the results of the event are reported on both DBSR fields. Setting DBCR0[IA12] =1 prohibits IAC1 and IAC2 from being used for exact address compares.

To enable Range 3:4, DBCR0[IA34] = 1 and DBCR0[IA3] or DBCR0[IA4] =1. An IAC event will be seen on the DBSR[IA n] field that corresponds to the enabled DBCR0[IA n] field. If DBCR0[IA3] and DBCR0[IA4] are enabled, the results of the event will be reported on both DBSR fields. Setting DBCR0[IA34] =1 prohibits IAC3 and IAC4 from being used for exact address compares.

Ranges can be defined as inclusive, as shown in the preceding examples, or exclusive, using DBCR0[IA12X] (corresponding to range 1:2) and DBCR0[IA34X] (corresponding to range 3:4), as follows:

- DBCR0[IA12] = 1: Range 1:2 = IAC1 ≤ range < IAC2.
- DBCR0[IA12X] = 1: Range 1:2 = Range low < IAC1 or IAC2 ≤ Range high
- DBCR0[IA34] = 1: Range 3:4 = IAC3 ≤ range < IAC4.
- DBCR0[IA34X] = 1: Range 3:4 = Range low < IAC3 or IAC4 ≤ Range high

Figure 8-7 shows the range selected in an inclusive IAC range address compare. Note that the address in IAC1 is considered part of the range, but the address in IAC2 is not, as shown in the preceding examples. The thick lines indicate that the indicated address is included in the compare results.




Figure 8-7. Inclusive IAC Range Address Compares

Figure 8-8 shows the range selected in an exclusive IAC range address compare. Note that the address in IAC1 is not considered part of the range, but the address in IAC2 is, along with the highest memory address, as shown in the preceding examples.




Figure 8-8. Exclusive IAC Range Address Compares

To toggle the range from inclusive to exclusive or from exclusive to inclusive on a IAC range debug event, DBCR0[IA12T] (corresponding to range 1:2) and DBCR0[IA34T] (corresponding to range 3:4) are used. If these fields are set, the DBCR0[IA12X] or DBCR0[IA34X] fields toggle on an IAC debug event, changing the defined range.

When a toggle is enabled (DBCR0[IA12T] for range 1:2 or DBCR0[IA34T] = 1 for range 3:4), and DBCR0[IDM] = 1, DBCR0[EDM] = 0, and MSR[DE] = 0, IAC range comparisons for the corresponding toggle field are disabled.

8.5.13 DAC Debug Event

This debug event occurs before execution of an instruction that accesses a data address that matches the contents of the specified DAC register. DBCR1[D1R, D2R, D1W, D2W] enable DAC debug events for address comparisons on DAC1 and DAC2 for read instructions, DAC2 for read instructions, DAC1 for write instructions, DAC2 for write instructions respectively. Loads are reads and stores are writes. DAC can be defined (DBCR1[D1R, D2R]) as an exact address comparison to one of the DACn registers or a range of addresses to compare defined by DAC1 and DAC2 registers.

8.5.13.1 DAC Exact Address Compare

In this mode, each DACn register specifies an exact address to compare. These registers are enabled by setting one or more of DBCR1[D1R, D2R, D1W, D2W] = 1, and disabling DAC range compare DBCR1[DA12X] = 0. The corresponding DBSR[DR1, DR2, DW1, DW2] field displays the results of a DAC debug event.

The address for a DAC is the effective address (EA) of a storage reference instruction. EAs are always generated within a single aligned word of memory. Unaligned load and store, strings, and multiples generate multiple EAs to be used in DAC comparisons.

Data address compare (DAC) debug events can be set to react to any byte in a larger block of memory, in addition to reacting to a byte address match. The DAC Compare Size fields (DBCR1[D1S, D2S]) allow DAC debug events to react to byte, halfword, word, or 8-word line address by ignoring a number of LSBs in the EA.

DAC 1 Size	
00 Compare all bits	Byte address
01 Ignore LSB (least significant bit)	Halfword address
10 Ignore two LSBs	Word address
11 Ignore five LSBs	Cache line (8-word) address

The user must determine how the addresses of interest are accessed, relative to byte, halfword, word, string, and unaligned storage instructions, and adjust the DAC compare size field appropriately to cover the addresses of interest.

For example, suppose that a DAC debug event should react to byte 3 of a word-aligned target. A DAC set for exact compare would not recognize a reference to that byte by load/store word or load/store halfword instructions, because the byte address is not the EA of such instructions. In such a case, the D1S field must be set for a wider capture range (for example, to ignore the two least significant bits (LSBs) if word operations to the misaligned byte are to be detected). The wider capture range may result in excess debug events (events that are within the specified capture range, but reflect byte operations in addition to the desired byte). Such excess debug events must be handled by software.

While load/store string instructions are inherently byte addressed the processor will generate EAs containing the largest portion of an aligned word address as possible. It may not be possible to DAC on a specific individual byte using load/store string instructions.

8.5.13.2 DAC Range Address Compare

In this mode, the pair of DAC1 and DAC2 registers are used to define a range of addresses to compare.

To enable DAC range, DBCR1[DA12] = 1 and one or more of DBCR1[D1R,D2R,D1W,D2W] = 1. The DAC event is seen on the DBSR[DR1,DR2,DW1,DW2] field that corresponds to the DBCR1[D1R,D2R,D1W,D2W] field that is enabled. For example, if DBCR1[D1R] and DBCR1[D2R] are enabled, the results of a DAC debug event are reported on DBSR[DR1, DR2]. Setting DBCR1[DA12] = 1 prohibits DAC1 and DAC2 from being used for exact address compares.

Ranges are defined to be inclusive or exclusive, using the DBCR1[DA12X], as follows:

DBCR1[DA12] = 1: Range = $DAC1 \leq \text{range} < DAC2$.

DBCR1[DA12X] = 1: Range = Range low < DAC1 or $DAC2 \leq \text{Range high}$.

Figure 8-9 shows the range selected in an inclusive DAC range address compare. Note that the address in DAC1 is considered part of the range, but the address in DAC2 is not, as shown in the

preceding examples. The thick lines indicate that the indicated address is included in the compare results.




Figure 8-9. Inclusive DAC Range Address Compares

Figure 8-10 shows the range selected in an exclusive DAC range address compare. Note that the address in DAC1 is not considered part of the range, but the address in DAC2 is, along with the highest memory address, as shown in the preceding examples.




Figure 8-10. Exclusive DAC Range Address Compares

The DAC Compare Size fields (DBCR1[D1S, D2S]) are not used by DAC range comparisons.

8.5.13.3 DAC Applied to Cache Instructions

Some cache instructions can cause DAC debug events. There are several special cases.

Table 8-2 summarizes possible DAC debug events by cache instruction:

Table 8-2. DAC Applied to Cache Instructions

Instruction	Possible DAC Debug Event	
	DAC-Read	DAC-Write
dcba	No	Yes
dcbf	No	Yes
dcbi	No	Yes
dcbst	No	Yes
dcbt	Yes	No
dcbz	No	Yes
dccci	No	No
dcread	No	No
dcbtst	Yes	No
icbi	Yes	No
icbt	Yes	No

Table 8-2. DAC Applied to Cache Instructions (continued)

Instruction	Possible DAC Debug Event	
	DAC-Read	DAC-Write
iccci	No	No
icread	No	No

Architecturally, the **dcbi** and **dcbz** instructions are “stores.” These instructions can change data, or cause the loss of data by invalidating a dirty line. Therefore, they can cause DAC-write debug events.

The **dccci** instruction can also be considered a “store” because it can change data by invalidating a dirty line. However, **dccci** is not address-specific; it affects an entire congruence class regardless of the operand address of the instruction. Because it is not address-specific, **dccci** does not cause DAC-write debug events.

Architecturally, the **dcbt**, **dcbtst**, **dcbf**, and **dcbst** instructions are “loads.” These instructions do not change data. Flushing or storing a cache line from the cache is not architecturally a “store” because a store had already updated the cache; the **dcbf** or **dcbst** instruction only updates the copy in main memory.

The **dcbt** and **dcbtst** instructions can cause DAC-read debug events regardless of cachability.

Although **dcbf** and **dcbst** are architecturally “loads,” these instructions can create DAC-write (but not DAC-read) debug events. In a debug environment, the fact that external memory is being written is the event of interest.

Even though **dcread** and **dccci** are not address-specific (they affect a congruence class regardless of the instruction operand address), and are considered “loads,” in the PPC405 they do not cause DAC debug events.

All ICU operations (**icbi**, **icbt**, **iccci**, and **icread**) are architecturally treated as “loads.” **icbi** and **icbt** cause DAC debug events. **iccci** and **icread** do not cause DAC debug events in the PPC405.

8.5.13.4 DAC Applied to String Instructions

An **stswx** instruction with a string length of 0 is a no-op. The **lswx** instruction with the string length equal to 0 does not alter the RT operand with undefined data, as allowed by the PowerPC Architecture. Neither **stswx** nor **lswx** with zero length causes a DAC debug event because storage is not accessed by these instructions.

8.5.14 Data Value Compare Debug Event

A data value compare (DVC) debug event can occur only after execution of a load or store instruction to an address that compares with the address in one of the DAC n registers and has a data value that matches the corresponding DVC n register. Therefore, a DVC debug event requires both the data address comparison and the data value comparison to be true. A DVC n debug event when enabled in the DBCR1 supercedes a DAC n debug event since the DVC n and the DAC n both use the same DAC n register.

DVC1 debug events are enabled by setting the appropriate DAC enable DBCR1[D1R,D1W] to cause an address comparison and by setting anybit combination in the DBCR1[DV1BE]. DVC2 debug events are enabled by setting the appropriate DAC enable DBCR1[D2R,D2W] to cause an address

comparison and by setting any bit combination in the DBCR1[DV1BE]. Each bit in DBCR1[DV1BE, DV2BE] corresponds to a byte in DVC1 and DVC2. Exact address compare and range address compare work the same for DVC as for a simple DAC.

DBSR[DR1] and DBSR[DW1] record status for DAC1 debug events. Which DBSR bit is set depends on the setting of DBCR1[D1R] and DBCR1[D1W]. If DBCR1[D1R] = 1, DBSR[DR1] = 1, assuming that a DVC event occurred. Similarly, if DBCR1[D1W] = 1, DBSR[DW1] = 1, assuming that a DVC event occurred.

Similarly, DBSR[DR2] and DBSR[DW2] record status for DAC2 debug events. Which DBSR bit is set depends on the setting of DBCR1[D2R] and DBCR1[D2W]. If DBCR1[D2R] = 1, DBSR[DR2] = 1, assuming that a DVC event occurred. Similarly, if DBCR1[D2W] = 1, DBSR[DW2] = 1, assuming that a DVC event occurred.

In the following example, a DVC1 event is enabled by setting DBCR1[D1R] = 1, DBCR1[D1W] = 1, DBCR1[DA12] = 0, and DBCR1[DV1BE] = 0000. When the data address and data value match the DAC1 and DVC1, a DVC1 event is recorded in DBSR[DR1] or DBSR[DW1], depending on whether the operation is a load (read) or a store (write). This example corresponds to the last line of Table 8-3.

In Table 8-3, *n* is 1 or 2, depending on whether the bits apply to DAC1, DAC2, DVC1, and DVC2 events. “Hold” indicates that the DBSR holds its value unless cleared by software. “RA” indicates that the operation is a read (load) and the data address compares (exact or range). “WA” indicates that the operation is a write (store) and the data address compares (exact or range). “RV” indicates that the operation is a read (load), the data address compares (exact or range), and the data value compares according to DBCR1[DVC*n*].

Table 8-3. Setting of DBSR Bits for DAC and DVC Events

DAC <i>n</i> Event	DVC <i>n</i> Enabled	DVC <i>n</i> Event	DBCR1			DBSR	
			[D <i>n</i> R]	[D <i>n</i> W]	[DA12]	[DR <i>n</i>]	[DW <i>n</i>]
0	—	—	—	—	—	Hold	Hold
—	—	—	0	0	—	Hold	Hold
1	0	—	0	1	—	Hold	WA
1	0	—	1	0	—	RA	Hold
1	0	—	1	1	—	RA	WA
1	1	0	—	—	—	Hold	Hold
1	1	1	0	1	—	Hold	WV
1	1	1	1	0	—	RV	Hold
1	1	1	1	1	—	RV	WV

The settings of DBCR1[DV1M] and DBCR1[DV2M] are more precisely defined in Table 8-5 and Table 8-6. (*n* enables the table to apply to DBCR1[DV1M, DV2M] and DBCR1[DV1BE, DV2BE]). DV*n*BE_{*m*} indicates bytes selected (or not selected) for comparison in DBCR1[DV*n*BE].

When DBCR1[DV*n*M] = 01, the comparison is an AND; all bytes must compare to the appropriate bytes of DVC1.

When DBCR1[DV*n*M] = 10, the comparison is an OR; at least one of the selected bytes must compare to the appropriate bytes of DVC1.

When DBCR1[DVnM] = 11, the comparison is an AND-OR (halfword) comparison. This is intended for use when DBCR1[DVnBE] is set to 0011, 0111, or 1111. Other values of DBCR1[DVnBE] can be compared, but the results are more easily understood using the AND and OR comparisons. In Table 8-4, “not” is \neg , AND is \wedge , and OR is \vee .

Table 8-4. Comparisons Based on DBCR1[DVnM]

DBCR1[DVnM] Setting	Operation	Comparison
00	—	Undefined
01	AND	$(\neg DVnBE_0 \vee (DVC1[\text{byte } 0] = \text{data}[\text{byte } 0])) \wedge$ $(\neg DVnBE_1 \vee (DVC1[\text{byte } 1] = \text{data}[\text{byte } 1])) \wedge$ $(\neg DVnBE_2 \vee (DVC1[\text{byte } 2] = \text{data}[\text{byte } 2])) \wedge$ $(\neg DVnBE_3 \vee (DVC1[\text{byte } 3] = \text{data}[\text{byte } 3]))$
10	OR	$(DVnBE_0 \wedge (DVC1[\text{byte } 0] = \text{data}[\text{byte } 0])) \vee$ $(DVnBE_1 \wedge (DVC1[\text{byte } 1] = \text{data}[\text{byte } 1])) \vee$ $(DVnBE_2 \wedge (DVC1[\text{byte } 2] = \text{data}[\text{byte } 2])) \vee$ $(DVnBE_3 \wedge (DVC1[\text{byte } 3] = \text{data}[\text{byte } 3]))$
11	AND-OR	$(DVnBE_0 \wedge (DVC1[\text{byte } 0] = \text{data}[\text{byte } 0])) \wedge$ $(DVnBE_1 \wedge (DVC1[\text{byte } 1] = \text{data}[\text{byte } 1])) \vee$ $(DVnBE_2 \wedge (DVC1[\text{byte } 2] = \text{data}[\text{byte } 2])) \wedge$ $(DVnBE_3 \wedge (DVC1[\text{byte } 3] = \text{data}[\text{byte } 1]))$

Table 8-5 illustrates comparisons for aligned DVC accesses, that is, words, halfwords, or bytes on naturally aligned boundaries (all byte accesses are aligned).

Table 8-5. Comparisons for Aligned DVC Accesses

Access	DBCR1[DVnBE] Setting	Value	Operation
Word	All	Word value	AND
Halfword (Low-Order)	All	Halfword value replicated	AND-OR
Halfword (High-Order)	All	Halfword value replicated	AND-OR
Byte	All	Byte value replicated	OR

For halfword accesses, the halfword value is replicated in the “empty “ halfword in the DVC register, for example, if the low-order halfword is to be compared, its value is stored in the low-order halfword and the high-order halfword of the register. Similarly, a byte value is replicated in each byte in the register.

Table 8-6 illustrates comparisons for misaligned DVC accesses. In the “DVC1” and “DVC2” columns, “x” indicates a don’t care.

Table 8-6. Comparisons for Misaligned DVC Accesses

Access	Operation	DVC1 (Hex)	DVC2 (Hex)	DBCR1[DV1BE] Setting	DBCR1[DV2BE] Setting	DBCR1[D2S] Setting
Word (Offset 1)	AND	xx112233	44xx xxxx	123	0	01
Word (Offset 2)	AND	xxxx1122	3344xxxx	23	01	10
Word (Offset 3)	AND	xxxxxx11	223344xx	3	012	10
Halfword (Offset 1)	AND	xx1122xx		12	12	10
Halfword (Offset 3)	AND	xxxxxx11	22xxxxxx	3	0	10

Note: Misaligned accesses stop the processor on the instruction causing the compare hit. The second part of an instruction is not performed if the first part of the compare hits.

8.5.15 Imprecise Debug Event

The imprecise debug event is not an independent debug event, but indicates that a debug event occurred while MSR[DE] = 0. This is useful in internal debug mode if a debug event occurs while in a critical interrupt handler. On return from interrupt, a debug interrupt occurs if MSR[DE] = 1. If DBSR[IDE] = 1, the debug event causing the interrupt occurred sometime earlier, not immediately after a debug event.

8.6 Debug Interface

The PPC405 core provides a and trace interfaces to support hardware and software test and debug. Typically, the JTAG interface connects to a debug port external to the PPC405; the debug port is typically connected to a JTAG connector on a processor board.

The trace interface connects to a trace port, also external to the PPC405, that is typically connected to a trace connector on the processor board.

8.6.1 IEEE 1149.1 Test Access Port (JTAG Debug Port)

The IEEE 1149.1 Test Access Port (TAP), commonly called the JTAG (Joint Test Action Group) debug port, is an architectural standard described in IEEE Std 1149.1–1990, *IEEE Standard Test Access Port and Boundary Scan Architecture*. The standard describes a method for accessing internal chip facilities using a four- or five-signal interface.

The JTAG debug port, originally designed to support scan-based board testing, is enhanced to support the attachment of debug tools. The enhancements, which comply with the IEEE 1149.1

specifications for vendor-specific extensions, are compatible with standard JTAG hardware for boundary-scan system testing.

- JTAG Signals** The JTAG debug port implements the four required JTAG signals: TCK, TMS, TDI, and TDO, and the optional $\overline{\text{TRST}}$ signal.
- JTAG Clock Requirements** The frequency of the TCK signal can range from DC to one-half of the internal chip clock frequency.
- JTAG Reset Requirements** The JTAG debug port logic is reset at the same time as a system reset. Upon receiving $\overline{\text{TRST}}$, the JTAG TAP controller returns to the Test-Logic Reset state.

8.7 JTAG Connector

A 16-pin male 2x8 header connector is suggested as the JTAG debug port connector. This connector definition matches the requirements of the RISCWatch debugger from IBM. The connector is shown in Figure 8-11 and the signals are shown in Table 8-7. The connector should be placed as close as possible to the chip to ensure signal integrity.

Note that position 14 does not contain a pin.




Figure 8-11. JTAG Connector Physical Layout (Top View)

Table 8-7. JTAG Connector Signals

Pin	I/O	Signal	Description
1	O	TDO	JTAG Test Data Out
2		No connect (NC)	Reserved
3	I	TDI ¹	JTAG Test Data In
4		$\overline{\text{TRST}}$	JTAG Reset
5		NC	Reserved
6		+POWER ²	Processor Power OK
7	I	TCK ³	JTAG Test Clock
8		NC	Reserved

Table 8-7. JTAG Connector Signals (continued)

Pin	I/O	Signal	Description
9	I	TMS ¹	JTAG Test Mode Select
10		NC	Reserved
11	I	HALT ³	Processor Halt
12		NC	Reserved
13		NC	Reserved
14		Key	The pin at this position should be removed.
15		NC	Reserved
16		GND	Ground

1. A 10K ohm pullup resistor should be connected to this signal to reduce chip power consumption. The pullup resistor is not required.
2. The +POWER signal, sourced from the target development board, indicates whether the processor is operating. This signal does not supply power to the RISCWatch hardware or to the processor. The active level on this signal can be +5V or +3.3V (note that the PPC405 core can have either +5V or +3.3V I/O, but the processor itself must be powered by +3.3V). A series resistor (1K ohm or less) should be used to provide short circuit current-limiting protection.
3. A 10K ohm pullup resistor must be connected to these signals to ensure proper chip operation when these inputs are not used.

8.7.1 JTAG Instructions

The JTAG debug port provides the standard *extest*, *idcode*, *sample/preload*, and *bypass* instructions and the optional *highz* and *clamp* instructions. Invalid instructions behave as the *bypass* instruction.

Table 8-8. JTAG Instructions

Instruction	Code	Comments
Extest	000	IEEE 1149.1 standard.
Intest	1111001	IEEE 1149.1 standard.
Sample/Preload	1111010	IEEE 1149.1 standard.
Private	xxxx100	Private instructions
Bypass	1111111	IEEE 1149.1 standard.

8.7.2 JTAG Boundary Scan

Boundary Scan Description Language (BSDL), IEEE 1149.1b-1994, is a supplement to IEEE 1149.1-1990 and IEEE 1149.1a-1993 *Standard Test Access Port and Boundary-Scan Architecture*. BSDL, a subset of the IEEE 1076-1993 Standard VHSIC Hardware Description Language (VHDL), allows a rigorous description of testability features in components which comply with the standard. BSDL is used by automated test pattern generation tools for package interconnect tests and by electronic design automation (EDA) tools for synthesized test logic and verification. BSDL supports

robust extensions that can be used for internal test generation and to write software for hardware debug and diagnostics.

The primary components of BSDL include the logical port description, the physical pin map, the instruction set, and the boundary register description.

The logical port description assigns symbolic names to the pins of a chip. Each pin has a logical type of in, out, inout, buffer, or linkage that defines the logical direction of signal flow.

The physical pin map correlates the logical ports of the chip to the physical pins of a specific package. A BSDL description can have several physical pin maps; each map is given a unique name.

Instruction set statements describe the bit patterns that must be shifted into the Instruction Register to place the chip in the various test modes defined by the standard. Instruction set statements also support descriptions of instructions that are unique to the chip.

The boundary register description lists each cell or shift stage of the Boundary Register. Each cell has a unique number: the cell numbered 0 is the closest to the Test Data Out (TDO) pin; the cell with the highest number is closest to the Test Data In (TDI) pin. Each cell contains additional information, including: cell type, logical port associated with the cell, logical function of the cell, safe value, control cell number, disable value, and result value.

8.8 Trace Port

The PPC405 core implements a trace status interface to support the tracing of code running in real-time. This interface enables the connection of an external trace tool, such as RISCWatch, and allows for user-extended trace functions. A software tool with trace capability, such as RISCWatch with RISCTrace, can use the data collected from this port to trace code running on the processor. The result is a trace of the code executed, including code executed out of the instruction cache if it was enabled. Information on trace capabilities, how trace works, and how to connect the external trace tool is available in *RISCWatch Debugger User's Guide*.

Chapter 9. Instruction Set

Descriptions of the PPC405 instructions follow. Each description contains the following elements:

- Instruction names (mnemonic and full)
- Instruction syntax
- Instruction format diagram
- Pseudocode description
- Prose description
- Registers altered
- Architecture notes identifying the associated PowerPC Architecture component

Where appropriate, instruction descriptions list invalid instruction forms and exceptions, and provide programming notes.

9.1 Instruction Set Portability

To support embedded real-time applications, the instruction sets of the PPC405 core and other IBM controllers implement the IBM PowerPC Embedded Environment, which is not part of the PowerPC Architecture defined in *The PowerPC Architecture: A Specification for a New Family of RISC Processors*.

Programs using these instructions are not portable to PowerPC implementations that do not implement the IBM PowerPC Embedded Environment.

The PPC405 core implements a number of implementation-specific instructions that are not part of the PowerPC Architecture or the IBM PowerPC Embedded Environment, which are listed in Table 9-1. In the table, the syntax “[o]” indicates that an instruction has an “o” form, which updates the XER[SO,OV] fields, and a “non-o” form. The syntax “[.]” indicates that an instruction has a “record” form, which updates CR[CR0], and a “non-record” form.

Table 9-1. Implementation-Specific Instructions

dccci	macchw[o][.]	mfdcr	nmacchw[o][.]	rfci
dcread	macchws[o][.]	mtdcr	nmacchws[o][.]	tlbre
iccci	macchwsu[o][.]	mulchw[.]	nmachhw[o][.]	tlbsx[.]
icread	macchwu[o][.]	mulchwu[.]	nmachhws[o][.]	tlbwe
	machhw[o][.]	mulhhw[.]	nmaclhw[o][.]	wrttee
	machhws[o][.]	mulhhwu[.]	nmaclhws[o][.]	wrtteei
	machhwsu[o][.]	mullhw[.]		
	machhwu[o][.]	mullhwu[.]		
	maclhw[o][.]			
	maclhws[o][.]			
	maclhwsu[o][.]			
	maclhwu[o][.]			

9.2 Instruction Formats

For more detailed information about instruction formats, including a summary of instruction field usage and instruction format diagrams for the PPC405 core, see “Instruction Formats” on page 9-2.

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. The 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 general purpose register selectors and immediate values, that may vary from execution to execution. The instruction format diagrams specify the operands in 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 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 and its result is architecturally undefined. Unless otherwise noted, the 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:

=	Assignment
\wedge	AND logical operator
\neg	NOT logical operator
\vee	OR logical operator
\oplus	Exclusive-OR (XOR) logical operator
+	Twos complement addition
-	Twos complement subtraction, unary minus
\times	Multiplication
\div	Division yielding a quotient
%	Remainder of an integer division; $(33 \% 32) = 1$.

	Concatenation
=, ≠	Equal, not equal relations
<, >	Signed comparison relations
⋚, ⋛	Unsigned comparison relations
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.
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.
leave	Leave innermost do loop or do loop specified in a leave statement.
n	A decimal number
0xn	A hexadecimal number
0bn	A binary number
FLD	An instruction or register field
FLD _b	A bit in a named instruction or register field
FLD _{b:b}	A range of bits in a named instruction or register field
FLD _{b,b, . . .}	A list of bits, by number or name, in a named instruction or register field
REG _b	A bit in a named register
REG _{b:b}	A range of bits in a named register
REG _{b,b, . . .}	A list of bits, by number or name, in a named register
REG[FLD]	A field in a named register
REG[FLD, FLD . . .]	A list of fields in a named register
REG[FLD:FLD]	A range of fields in a named register
GPR(<i>r</i>)	General Purpose Register (GPR) <i>r</i> , where $0 \leq r \leq 31$.
(GPR(<i>r</i>))	The contents of GPR <i>r</i> , where $0 \leq r \leq 31$.
DCR(DCRN)	A Device Control Register (DCR) specified by the DCRF field in an mf dcr or mt dcr instruction
SPR(SPRN)	An SPR specified by the SPRF field in an mf spr or mt spr instruction
TBR(TBRN)	A Time Base Register (TBR) specified by the TBRF field in an mftb instruction
GPRs	RA, RB, . . .
(Rx)	The contents of a GPR, where <i>x</i> is A, B, S, or T
(RA 0)	The contents of the register RA or 0, if the RA field is 0.
c _{0:3}	A four-bit object used to store condition results in compare instructions.
"b	The bit or bit value <i>b</i> is replicated <i>n</i> times.

xx	Bit positions which are don't-cares.
CEIL(x)	Least integer $\geq x$.
EXTS(x)	The result of extending x on the left with sign bits.
PC	Program counter.
RESERVE	Reserve bit; indicates whether a process has reserved a block of storage.
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.
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.
MS(addr, n)	The number of bytes represented by n at the location in main storage represented by $addr$.
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.
EA _b	A bit in an effective address.
EA _{b:b}	A range of bits in an effective address.
ROTL((RS),n)	Rotate left; the contents of RS are shifted left the number of bits specified by n .
MASK(MB,ME)	Mask having 1s in positions MB through ME (wrapping if MB > ME) and 0s elsewhere.
instruction(EA)	An instruction operating on a data or instruction cache block associated with an EA.

9.3.1 Operator Precedence

Table 9-2 lists the pseudocode operators and their associativity in descending order of precedence:

Table 9-2. Operator Precedence

Operators	Associativity
REG _D , REG[FLD], function evaluation	Left to right
ⁿ b	Right to left
¬, − (unary minus)	Right to left
×, ÷	Left to right
+, −	Left to right
	Left to right
=, ≠, <, >, < ^u , > ^u	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). Other registers are changed, with the details of the change not included in the instruction description. This category frequently includes the Condition Register (CR) and the Fixed-point Exception Register (XER). For discussion of the CR, see “Condition Register (CR)” on page 2-10. For discussion of XER, see “Fixed Point Exception Register (XER)” on page 2-7.


9.5 Alphabetical Instruction Listing

The following pages list the instructions available in the PPC405 core in alphabetical order.

add

Add

add	RT, RA, RB	OE=0, Rc=0
add.	RT, RA, RB	OE=0, Rc=1
addo	RT, RA, RB	OE=1, Rc=0
addo.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

addc

Add Carrying

addc	RT, RA, RB	OE=0, Rc=0
addc.	RT, RA, RB	OE=0, Rc=1
addco	RT, RA, RB	OE=1, Rc=0
addco.	RT, RA, RB	OE=1, Rc=1



```
(RT) ← (RA) + (RB)
if (RA) + (RB)  $\geq$   $2^{32} - 1$  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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

adde

Add Extended

adde	RT, RA, RB	OE=0, Rc=0
adde.	RT, RA, RB	OE=0, Rc=1
addeo	RT, RA, RB	OE=1, Rc=0
addeo.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-3. Extended Mnemonics for addi

Mnemonic	Operands	Function	Other Registers Altered
la	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> addi RT,RA,D	
li	RT, IM	Load immediate. (RT) ← EXTS(IM) <i>Extended mnemonic for</i> addi RT,0,IM	
subi	RT, RA, IM	Subtract EXTS(IM) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> addi RT,RA,-IM	

addic

Add Immediate Carrying

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]

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-4. Extended Mnemonics for addic

Mnemonic	Operands	Function	Other Registers Altered
subic	RT, RA, IM	Subtract EXTS(IM) from (RA) Place result in RT; place carry-out in XER[CA]. <i>Extended mnemonic for</i> addic RT,RA,-IM	

addic.

Add Immediate Carrying and Record

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]_{LT, GT, EQ, SO}

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.**

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-5. Extended Mnemonics for addic.

Mnemonic	Operands	Function	Other Registers Altered
subic.	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

addis RT, RA, IM



$$(RT) \leftarrow (RA|0) + (IM \parallel 160)$$

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
```

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-6. Extended Mnemonics for addis

Mnemonic	Operands	Function	Other Registers Altered
lis	RT, IM	Load immediate shifted. $(RT) \leftarrow (IM \parallel 160)$ <i>Extended mnemonic for</i> addis RT,0,IM	
subis	RT, RA, IM	Subtract $(IM \parallel 160)$ from $(RA 0)$. Place result in RT. <i>Extended mnemonic for</i> addis RT,RA,-IM	

addme

Add to Minus One Extended

addme	RT, RA	OE=0, Rc=0
addme.	RT, RA	OE=0, Rc=1
addmeo	RT, RA	OE=1, Rc=0
addmeo.	RT, RA	OE=1, Rc=1



```
(RT) ← (RA) + XER[CA] + (-1)
if (RA) + XER[CA] + 0xFFFF FFFF  $\geq$   $2^{32} - 1$  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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

addze

Add to Zero Extended

addze	RT, RA	OE=0, Rc=0
addze.	RT, RA	OE=0, Rc=1
addzeo	RT, RA	OE=1, Rc=0
addzeo.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

and

AND

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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

andc

AND with Complement

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]_{LT,GT,EQ,SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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]_{LT, GT, EQ, SO}

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..**


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

andis.

AND Immediate Shifted

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]_{LT, GT, EQ, SO}

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 R_c field. The other instructions are **addic.** and **andi..**

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

b	target	AA=0, LK=0
ba	target	AA=1, LK=0
bl	target	AA=0, LK=1
bla	target	AA=1, LK=1



```

If AA = 1 then
  LI ← target6:29
  NIA ← EXTS(LI || 200)
else
  LI ← (target - CIA)6:29
  NIA ← CIA + EXTS(LI || 200)
if LK = 1 then
  (LR) ← CIA + 4
PC ← NIA

```

The next instruction address (NIA) is the effective address of the branch. 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 also the current instruction address (CIA). If the AA field contains 1, the base address is 0.

Program flow is transferred to the NIA.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

Registers Altered

- LR if LK contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

bc

Branch Conditional

bc	BO, BI, target	AA=0, LK=0
bca	BO, BI, target	AA=1, LK=0
bcl	BO, BI, target	AA=0, LK=1
bcla	BO, BI, target	AA=1, LK=1

16	BO	BI	BD	AA	LK
0	6	11	16	30	31

```
if  $BO_2 = 0$  then
    CTR  $\leftarrow$  CTR - 1
if  $(BO_2 = 1 \vee ((CTR = 0) = BO_3)) \wedge (BO_0 = 1 \vee (CR_{BI} = BO_1))$  then
    if AA = 1 then
        BD  $\leftarrow$  target16:29
        NIA  $\leftarrow$  EXTS(BD || 200)
    else
        BD  $\leftarrow$  (target - CIA)16:29
        NIA  $\leftarrow$  CIA + EXTS(BD || 200)
    else
        NIA  $\leftarrow$  CIA + 4
    if LK = 1 then
        (LR)  $\leftarrow$  CIA + 4
    PC  $\leftarrow$  NIA
```

If bit 2 of the BO field contains 0, the CTR decrements.

The BI field specifies a bit in the CR to be used as the condition of the branch.

The next instruction address (NIA) is the effective address of the branch. 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 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 also the current instruction address (CIA). If the AA field contains 1, the base address is 0.

The BO field controls options that determine when program flow is transferred to the NIA. The BO field also controls branch prediction, a performance-improvement feature. See “Branch Prediction” on page 2-26 for a complete discussion.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

Registers Altered

- CTR if BO_2 contains 0
- LR if LK contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla

Mnemonic	Operands	Function	Other Registers Altered
bdnz	target	Decrement CTR; branch if CTR \neq 0. <i>Extended mnemonic for</i> bc 16,0,target	
bdnza		<i>Extended mnemonic for</i> bca 16,0,target	
bdnzl		<i>Extended mnemonic for</i> bcl 16,0,target	(LR) \leftarrow CIA + 4.
bdnzla		<i>Extended mnemonic for</i> bcla 16,0,target	(LR) \leftarrow CIA + 4.
bdnzf	cr_bit, target	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 0,cr_bit,target	
bdnzfa		<i>Extended mnemonic for</i> bca 0,cr_bit,target	
bdnzfl		<i>Extended mnemonic for</i> bcl 0,cr_bit,target	(LR) \leftarrow CIA + 4.
bdnzfla		<i>Extended mnemonic for</i> bcla 0,cr_bit,target	(LR) \leftarrow CIA + 4.
bdnzf	cr_bit, target	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 8,cr_bit,target	
bdnzfa		<i>Extended mnemonic for</i> bca 8,cr_bit,target	
bdnzfl		<i>Extended mnemonic for</i> bcl 8,cr_bit,target	(LR) \leftarrow CIA + 4.
bdnzfla		<i>Extended mnemonic for</i> bcla 8,cr_bit,target	(LR) \leftarrow CIA + 4.
bdz	target	Decrement CTR; branch if CTR = 0. <i>Extended mnemonic for</i> bc 18,0,target	
bdza		<i>Extended mnemonic for</i> bca 18,0,target	
bdzl		<i>Extended mnemonic for</i> bcl 18,0,target	(LR) \leftarrow CIA + 4.
bdzla		<i>Extended mnemonic for</i> bcla 18,0,target	(LR) \leftarrow CIA + 4.

bc

Branch Conditional

Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
bdzf	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 2,cr_bit,target	
bdzfa		<i>Extended mnemonic for</i> bca 2,cr_bit,target	
bdzfl		<i>Extended mnemonic for</i> bcl 2,cr_bit,target	(LR) ← CIA + 4.
bdzfla		<i>Extended mnemonic for</i> bcla 2,cr_bit,target	(LR) ← CIA + 4.
bdzt	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 10,cr_bit,target	
bdzta		<i>Extended mnemonic for</i> bca 10,cr_bit,target	
bdztl		<i>Extended mnemonic for</i> bcl 10,cr_bit,target	(LR) ← CIA + 4.
bdztla		<i>Extended mnemonic for</i> bcla 10,cr_bit,target	(LR) ← CIA + 4.
beq	[cr_field,] target	Branch if equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+2,target	
beqa		<i>Extended mnemonic for</i> bca 12,4*cr_field+2,target	
beql		<i>Extended mnemonic for</i> bcl 12,4*cr_field+2,target	(LR) ← CIA + 4.
beqla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+2,target	(LR) ← CIA + 4.
bf	cr_bit, target	Branch if CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 4,cr_bit,target	
bfa		<i>Extended mnemonic for</i> bca 4,cr_bit,target	
bfl		<i>Extended mnemonic for</i> bcl 4,cr_bit,target	LR
bfla		<i>Extended mnemonic for</i> bcla 4,cr_bit,target	LR

Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
bge	[cr_field,] target	Branch if greater than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target	
bgea		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target	
bgel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	LR
bgea		<i>Extended mnemonic for</i> bcla 4,4*cr_field+0,target	LR
bgt	[cr_field,] target	Branch if greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+1,target	
bgt		<i>Extended mnemonic for</i> bca 12,4*cr_field+1,target	
bgtl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+1,target	LR
bgtla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+1,target	LR
ble	[cr_field,] target	Branch if less than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target	
blea		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target	
blel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	LR
blela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	LR
blt	[cr_field,] target	Branch if less than Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+0,target	
blt		<i>Extended mnemonic for</i> bca 12,4*cr_field+0,target	
bltl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+0,target	(LR) ← CIA + 4.
bltla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+0,target	(LR) ← CIA + 4.

bc

Branch Conditional

Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
bne	[cr_field,] target	Branch if not equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+2,target	
bnea		<i>Extended mnemonic for</i> bca 4,4*cr_field+2,target	
bnel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+2,target	(LR) ← CIA + 4.
bnela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+2,target	(LR) ← CIA + 4.
bng	[cr_field,] target	Branch if not greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target	
bnga		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target	
bngl		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	(LR) ← CIA + 4.
bngla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	(LR) ← CIA + 4.
bnl	[cr_field,] target	Branch if not less than; use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target	
bnla		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target	
bnll		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	(LR) ← CIA + 4.
bnlla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+0,target	(LR) ← CIA + 4.
bns	[cr_field,] target	Branch if not summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target	
bnsa		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target	
bnsl		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.
bnsla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.


Table 9-7. Extended Mnemonics for bc, bca, bcl, bcla (continued)

Mnemonic	Operands	Function	Other Registers Altered
bnu	[cr_field,] target	Branch if not unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target	
bnu		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target	
bnul		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.
bnula		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.
bso	[cr_field,] target	Branch if summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target	
bso		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target	
bsol		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) ← CIA + 4.
bsola		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) ← CIA + 4.
bt	cr_bit, target	Branch if CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 12,cr_bit,target	
bt		<i>Extended mnemonic for</i> bca 12,cr_bit,target	
btl		<i>Extended mnemonic for</i> bcl 12,cr_bit,target	(LR) ← CIA + 4.
btla		<i>Extended mnemonic for</i> bcla 12,cr_bit,target	(LR) ← CIA + 4.
bun	[cr_field], target	Branch if unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target	
bun		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target	
bunl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) ← CIA + 4.
bunla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) ← CIA + 4.

bcctr

Branch Conditional to Count Register

bcctr BO, BI LK=0
bcctrl BO, BI LK=1



```
if  $BO_2 = 0$  then
     $CTR \leftarrow CTR - 1$ 
if  $(BO_2 = 1 \vee ((CTR = 0) = BO_3)) \wedge (BO_0 = 1 \vee (CR_{BI} = BO_1))$  then
     $NIA \leftarrow CTR_{0:29} \parallel 20$ 
else
     $NIA \leftarrow CIA + 4$ 
if LK = 1 then
     $(LR) \leftarrow CIA + 4$ 
 $PC \leftarrow NIA$ 
```

The BI field specifies a bit in the CR to be used as the condition of the branch.

The next instruction address (NIA) is the target address of the branch. The NIA is formed by concatenating the 30 most significant bits of the CTR with two 0-bits on the right.

The BO field controls options that determine when program flow is transferred to the NIA. The BO field also controls branch prediction, a performance-improvement feature. See “Branch Prediction” on page 2-26 for a complete discussion.

If the LK field contains 1, then $(CIA + 4)$ is placed into the LR.

Registers Altered

- CTR if BO_2 contains 0
- LR if LK contains 1

Invalid Instruction Forms

- Reserved fields
- If bit 2 of the BO field contains 0, the instruction form is invalid, but the pseudocode applies. If the branch condition is true, the branch is taken; the NIA is the contents of the CTR after it is decremented.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-8. Extended Mnemonics for bcctr, bcctrl

Mnemonic	Operands	Function	Other Registers Altered
bctr		Branch unconditionally to address in CTR. <i>Extended mnemonic for</i> bcctr 20,0	
bctrl		<i>Extended mnemonic for</i> bcctrl 20,0	(LR) ← CIA + 4.
beqctr	[cr_field]	Branch, if equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+2	
beqctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+2	(LR) ← CIA + 4.
bfctr	cr_bit	Branch, if CR _{cr_bit} = 0, to address in CTR. <i>Extended mnemonic for</i> bcctr 4,cr_bit	
bfctrl		<i>Extended mnemonic for</i> bcctrl 4,cr_bit	(LR) ← CIA + 4.
bgectr	[cr_field]	Branch, if greater than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0	
bgectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.
bgtctr	[cr_field]	Branch, if greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+1	
bgtctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+1	(LR) ← CIA + 4.
blectr	[cr_field]	Branch, if less than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1	
blectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.
bltctr	[cr_field]	Branch, if less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+0	
bltctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+0	(LR) ← CIA + 4.

bcctr

Branch Conditional to Count Register

Table 9-8. Extended Mnemonics for bcctr, bcctrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
bnectr	[cr_field]	Branch, if not equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+2	
bnectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+2	(LR) ← CIA + 4.
bngctr	[cr_field]	Branch, if not greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1	
bngctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.
bnlctr	[cr_field]	Branch, if not less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0	
bnlctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.
bnsctr	[cr_field]	Branch, if not summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3	
bnsctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.
bnuctr	[cr_field]	Branch, if not unordered, to address in CTR; use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3	
bnuctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.
bsocctr	[cr_field]	Branch, if summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3	
bsocctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) ← CIA + 4.
btctr	cr_bit	Branch if CR _{cr_bit} = 1 to address in CTR. <i>Extended mnemonic for</i> bcctr 12,cr_bit	
btctrl		<i>Extended mnemonic for</i> bcctrl 12,cr_bit	(LR) ← CIA + 4.

Table 9-8. Extended Mnemonics for bcctr, bcctrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
bunctr	[cr_field]	Branch if unordered to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3	
bunctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) ← CIA + 4.

bclr

Branch Conditional to Link Register

bclr BO, BI LK = 0
bclrl 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 bit 2 of the BO field contains 0, the CTR is decremented.

The BI field specifies a bit in the CR to be used as the condition of the branch.

The next instruction address (NIA) is the target address of the branch. The NIA is formed by concatenating the 30 most significant bits of the LR with two 0-bits on the right.

The BO field controls options that determine when program flow is transferred to the NIA. The BO field also controls branch prediction, a performance-improvement feature. See “Branch Prediction” on page 2-26 for a complete discussion.

If the LK field contains 1, then (CIA + 4) is placed into the LR.

Registers Altered

- CTR if BO₂ contains 0
- LR if LK contains 1

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-9. Extended Mnemonics for bclr, bclrl

Mnemonic	Operands	Function	Other Registers Altered
blr		Branch unconditionally to address in LR. <i>Extended mnemonic for</i> bclr 20,0	
bclrl		<i>Extended mnemonic for</i> bclrl 20,0	(LR) ← CIA + 4.

Table 9-9. Extended Mnemonics for bclr, bclrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
bdnzlr		Decrement CTR. Branch if CTR \neq 0 to address in LR. <i>Extended mnemonic for</i> bclr 16,0	
bdnzlrl		<i>Extended mnemonic for</i> bclrl 16,0	(LR) \leftarrow CIA + 4.
bdnzflr	cr_bit	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 0,cr_bit	
bdnzflrl		<i>Extended mnemonic for</i> bclrl 0,cr_bit	(LR) \leftarrow CIA + 4.
bdnztlr	cr_bit	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 1 to address in LR. <i>Extended mnemonic for</i> bclr 8,cr_bit	
bdnztlrl		<i>Extended mnemonic for</i> bclrl 8,cr_bit	(LR) \leftarrow CIA + 4.
bdzlr		Decrement CTR. Branch if CTR = 0 to address in LR. <i>Extended mnemonic for</i> bclr 18,0	
bdzlrl		<i>Extended mnemonic for</i> bclrl 18,0	(LR) \leftarrow CIA + 4.
bdzflr	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 2,cr_bit	
bdzflrl		<i>Extended mnemonic for</i> bclrl 2,cr_bit	(LR) \leftarrow CIA + 4.
bdztlr	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 1 to address in LR. <i>Extended mnemonic for</i> bclr 10,cr_bit	
bdztlrl		<i>Extended mnemonic for</i> bclrl 10,cr_bit	(LR) \leftarrow CIA + 4.
beqlr	[cr_field]	Branch if equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+2	
beqlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+2	(LR) \leftarrow CIA + 4.

bclr

Branch Conditional to Link Register

Table 9-9. Extended Mnemonics for bclr, bclrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
bflr	cr_bit	Branch if $CR_{cr_bit} = 0$ to address in LR. <i>Extended mnemonic for</i> bclr 4,cr_bit	
bflrl		<i>Extended mnemonic for</i> bclrl 4,cr_bit	(LR) \leftarrow CIA + 4.
bgehr	[cr_field]	Branch, if greater than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0	
bgehrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) \leftarrow CIA + 4.
bgtlr	[cr_field]	Branch, if greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+1	
bgtlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+1	(LR) \leftarrow CIA + 4.
blehr	[cr_field]	Branch, if less than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1	
blehrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) \leftarrow CIA + 4.
bltlr	[cr_field]	Branch, if less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+0	
bltlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+0	(LR) \leftarrow CIA + 4.
bnelr	[cr_field]	Branch, if not equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+2	
bnelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+2	(LR) \leftarrow CIA + 4.
bnglr	[cr_field]	Branch, if not greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1	
bnglrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) \leftarrow CIA + 4.


Table 9-9. Extended Mnemonics for bclr, bclrl (continued)

Mnemonic	Operands	Function	Other Registers Altered
bnllr	[cr_field]	Branch, if not less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0	
bnllrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) ← CIA + 4.
bnslr	[cr_field]	Branch if not summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3	
bnsrlr		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.
bnulr	[cr_field]	Branch if not unordered to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3	
bnulrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.
bsolr	[cr_field]	Branch if summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3	
bsolrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) ← CIA + 4.
btlr	cr_bit	Branch if CR _{cr_bit} = 1 to address in LR. <i>Extended mnemonic for</i> bclr 12,cr_bit	
btlrl		<i>Extended mnemonic for</i> bclrl 12,cr_bit	(LR) ← CIA + 4.
bunlr	[cr_field]	Branch if unordered to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3	
bunlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) ← CIA + 4.

cmp

Compare

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

The PowerPC Architecture defines this instruction as **cmp BF,L,RA,RB**, where L selects operand size for 64-bit PowerPC implementations. For all 32-bit PowerPC implementations, L = 0 is required (L = 1 is an invalid form); hence for PPC405 core, use of the extended mnemonic **cmpw BF,RA,RB** is recommended.


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-10. Extended Mnemonics for cmp

Mnemonic	Operands	Function	Other Registers Altered
cmpw	[BF,] RA, RB	Compare Word; use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmp BF,0,RA,RB	

cmpi BF, 0, RA, IM



```

c0:3 ← 40
if (RA) < EXTS(IM) then c0 ← 1
if (RA) > EXTS(IM) then c1 ← 1
if (RA) = EXTS(IM) then c2 ← 1
c3 ← XER[SO]
n ← BF
CR[CRn] ← c0:3
    
```

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[CRn] where n is specified by the BF field

Invalid Instruction Forms

- Reserved fields

Programming Note

The PowerPC Architecture defines this instruction as **cmpi BF,L,RA,IM**, where L selects operand size for 64-bit PowerPC implementations. For all 32-bit PowerPC implementations, L = 0 is required (L = 1 is an invalid form); hence for the PPC405 core, use of the extended mnemonic **cmpwi BF,RA,IM** is recommended.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-11. Extended Mnemonics for cmpi

Mnemonic	Operands	Function	Other Registers Altered
cmpwi	[BF,] RA, IM	Compare Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmpi BF,0,RA,IM	

cmpl

Compare Logical

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

The PowerPC Architecture defines this instruction as **cmpl BF,L,RA,RB**, where L selects operand size for 64-bit PowerPC implementations. For all 32-bit PowerPC implementations, L = 0 is required (L = 1 is an invalid form); hence for PPC405 core, use of the extended mnemonic **cmplw BF,RA,RB** is recommended.


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-12. Extended Mnemonics for cmpl

Mnemonic	Operands	Function	Other Registers Altered
cmplw	[BF,] RA, RB	Compare Logical Word. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmpl BF,0,RA,RB	

cmpli BF, 0, RA, IM



```

c0:3 ← 40
if (RA) u < (160 || IM) then c0 ← 1
if (RA) u > (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[CR_n] where *n* is specified by the BF field

Invalid Instruction Forms

- Reserved fields

Programming Note

The PowerPC Architecture defines this instruction as **cmpli BF,L,RA,IM**, where L selects operand size for 64-bit PowerPC implementations. For all 32-bit PowerPC implementations, L = 0 is required (L = 1 is an invalid form); hence for the PPC405 core, use of the extended mnemonic **cmplwi BF,RA,IM** is recommended.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-13. Extended Mnemonics for cmpli

Mnemonic	Operands	Function	Other Registers Changed
cmplwi	[BF,] RA, IM	Compare Logical Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for cmpli BF,0,RA,IM</i>	

cntlzw

Count Leading Zeros Word

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]_{LT, GT, EQ, SO} if Rc contains 1

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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.

Registers Altered

- CR

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

crandc

Condition Register AND with Complement

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.

Registers Altered

- CR


Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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.

Registers Altered

- CR

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-14. Extended Mnemonics for creqv

Mnemonic	Operands	Function	Other Registers Altered
crset	bx	CR set. <i>Extended mnemonic for creqv bx,bx,bx</i>	

crnand

Condition Register NAND

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.

Registers Altered

- CR


Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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.

Registers Altered

- CR

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.


Table 9-15. Extended Mnemonics for crnor

Mnemonic	Operands	Function	Other Registers Altered
crnot	bx, by	CR not. <i>Extended mnemonic for crnor bx,by,by</i>	

cror

Condition Register OR

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.

Registers Altered

- CR

Invalid Instruction Forms


- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-16. Extended Mnemonics for cror

Mnemonic	Operands	Function	Other Registers Altered
crmove	bx, by	CR move. <i>Extended mnemonic for cror bx,by,by</i>	

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.

Registers Altered

- CR

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

CRXOR

Condition Register XOR

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.

Registers Altered

- CR

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-17. Extended Mnemonics for crxor

Mnemonic	Operands	Function	Other Registers Altered
crclr	bx	Condition register clear. <i>Extended mnemonic for crxor bx,bx,bx</i>	

dcba RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$DCBA(EA)$$

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 EA is marked as cachable and non-write-through, the data in the cache block is architecturally undefined. For the PPC405 core, the cache data block is set to 0.

If the data block at the EA is not in the data cache and the EA is marked as cachable and not marked as write-through, a cache block is established and set to an architecturally-undefined value. Note that no data is read from main storage, as described in the programming note.

If the data block at the EA is marked as non-cachable, a no-op occurs.

If the data block at the EA is in the data cache and marked as write-through, architecturally the data in the cache block can be left unmodified. Alternatively, the data block at the EA can be undefined in the data cache and in main storage. For the PPC405 core, a no-op occurs.

If the data block at the EA is not in the data cache and marked as write-through, architecturally the instruction can establish a cache block and set the block to 0, or a no-op can occur. For the PPC405 core, a no-op 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 **dcba** can establish an address in the data cache without copying the contents of that address from main storage, the address established can be invalid with respect to main storage. 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 machine check exception could occur under these circumstances.

dcba provides a hint that a block of storage will soon be stored to or no longer needed; there is no need to retain the data in the block. Establishing the line in the cache, without reading from main storage, improves performance.

dcba

Data Cache Block Allocate

Exceptions

This instruction is considered a “store” with respect to data storage exceptions. However, this instruction does not cause data storage exceptions or data TLB-miss exceptions. If conditions occur that would otherwise cause such exceptions, **dcba** is treated as a no-op.

This instruction is considered a “store” with respect to data address compare (DAC) debug exceptions. See “Data Storage Interrupt” on page 5-16.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

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 EA is marked as cachable.

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 5-16.

This instruction is considered a “store” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

dcbi

Data Cache Block Invalidate

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 EA is marked as cachable. 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 5-16.

This instruction is considered a “store” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

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 EA is marked as cachable.

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 5-16.

This instruction is considered a “store” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

dcbt

Data Cache Block Touch

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 EA is marked as cachable, 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 EA is marked as non-cachable, no operation is performed.

This instruction is not allowed to cause data storage exceptions or data TLB miss exceptions. If execution of the instruction would cause such an exception, then no operation is performed, and no 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

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 5-16.

This instruction is considered a “load” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

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 EA address is marked as cachable, the data block is loaded into the data cache.

If the EA is marked as non-cachable, or if the data block at the EA is in the data cache, no operation is performed.

This instruction is not allowed to cause data storage exceptions or data TLB miss exceptions. If execution of the instruction would cause such an exception, then no operation is performed, and no 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

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** brings data into the cache in “Exclusive” mode, which allows the program to alter the cached data. “Exclusive” mode is part of the MESI protocol for multi-processor systems, and is not implemented. 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 5-16.

This instruction is considered a “load” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

dcbz

Data Cache Block Set to Zero

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 EA is marked as cachable and non-write-through, the data in the cache block is set to 0.

If the data block at the EA is not in the data cache and the EA is marked as cachable and non-write-through, a cache block is established and set to 0. Note that nothing is read from main storage, as described in the programming note.

If the data block at the EA is marked as either write-through or as non-cachable, 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 machine check exception could occur under these circumstances.

If **dcbz** is attempted to an EA which is marked as non-cachable, the software alignment exception handler should emulate the instruction by storing zeros to the block in main storage. If a data block corresponding to the EA exists in the cache, but the EA is non-cachable, stores (including **dcbz**) to that address are considered programming errors (the cache block should previously have been flushed).

If the EA is marked as write-through, the software alignment exception handler should emulate the instruction by storing zeros to the block in main storage. An EA that is marked as write-through required should also be marked as cachable; when **dcbz** is attempted to such an address, the alignment exception handler should maintain coherency of cache and memory.

Exceptions

An alignment exception occurs if the EA is marked as non-cachable or as write-through.

This instruction is considered a “store” with respect to data storage exceptions. See “Data Storage Interrupt” on page 5-16.

This instruction is considered a “store” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

dccci

Data Cache Congruence Class Invalidate

dccci RA, RB



$EA \leftarrow (RA|0) + (RB)$
DCCCI(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.

Both cache lines in the congruence class specified by $EA_{18:26}$ are invalidated, whether or not they match the EA. If modified data existed in the cache congruence class before the operation of this instruction, that data is lost.

The operation specified by this instruction is performed whether or not the EA is marked as cachable.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Registers Altered

- None

Invalid Instruction Forms

- Reserved fields

Programming Note

Execution of this instruction is privileged.

This instruction is intended for use in the power-on reset routine to invalidate the entire data cache tag array before enabling the data cache. A series of **dccci** instruction should be executed, one for each congruence class. Cachability can then be enabled.

Exceptions

See “Access Protection for Cache Control Instructions” on page 7-16.

The execution of an **dccci** instruction can cause a data TLB miss exception, at the specified EA, regardless of the non-specific intent of that EA.

This instruction does not cause data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.

Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

dcread RT, RA, RB

31	RT	RA	RB	486	
0	6	11	16	21	31

$EA \leftarrow (RA|0) + (RB)$
 if $((CCR0[CIS] = 0) \wedge (CCR0[CWS] = 0))$ then $(RT) \leftarrow$ (d-cache data, way A)
 if $((CCR0[CIS] = 0) \wedge (CCR0[CWS] = 1))$ then $(RT) \leftarrow$ (d-cache data, way B)
 if $((CCR0[CIS] = 1) \wedge (CCR0[CWS] = 0))$ then $(RT) \leftarrow$ (d-cache tag, way A)
 if $((CCR0[CIS] = 1) \wedge (CCR0[CWS] = 1))$ then $(RT) \leftarrow$ (d-cache tag, way B)

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.

This instruction is a debugging tool for reading the data cache entries for the congruence class specified by $EA_{18:26}$, unless no cache array is present. The cache information is read into register RT.

If $CCR0[CIS] = 0$, the information is a word of data cache array data from the addressed congruence class. The word is specified by $EA_{27:29}$. If $EA_{30:31}$ are not 00, an alignment exception occurs. If $CCR0[CWS] = 0$, the data is from the A-way; otherwise, the data is from the B-way.

If $CCR0[CIS] = 1$, the information is a cache tag from the addressed congruence class. If $CCR0[CWS] = 0$, the tag is from the A-way; otherwise the tag is from the B-way.

Data cache tag information is placed into register RT as shown:

0:19	TAG	Cache Tag
20:25		Reserved
26	D	Cache Line Dirty 0 Not dirty 1 Dirty
27	V	Cache Line Valid 0 Not valid 1 Valid
28:30		Reserved
31	LRU	Least Recently Used (LRU) 0 A-way LRU 1 B-way LRU

If instruction bit 31 contains 1, the contents of $CR[CR0]$ are undefined.

Registers Altered

- RT

Invalid Instruction Forms

- Reserved fields

dcread

Data Cache Read

Programming Note

Execution of this instruction is privileged.

Exceptions

If EA is not word-aligned, an alignment exception occurs.

This instruction is considered a “load” with respect to data storage exceptions, but cannot cause a data storage exception. See “Access Protection for Cache Control Instructions” on page 7-16.

The execution of an **dcread** instruction can cause a data TLB miss exception, at the specified EA, regardless of the non-specific intent of that effective address.

This instruction is considered a “load” with respect to data address compare (DAC) debug exceptions. See “Debug Interrupt” on page 5-26.

Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

divw	RT, RA, RB	OE=0, Rc=0
divw.	RT, RA, RB	OE=0, Rc=1
divwo	RT, RA, RB	OE=1, Rc=0
divwo.	RT, RA, RB	OE=1, Rc=1

31	RT	RA	RB	OE	491	Rc
0	6	11	16	21 22		31

$$(RT) \leftarrow (RA) \div (RB)$$

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 Rc field also contains 1, the contents of CR[CR0]_{LT, GT, EQ} are undefined. Either invalid division operation sets XER[OV, SO] to 1 if the OE field contains 1.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} 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:

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.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

divwu

Divide Word Unsigned

divwu	RT, RA, RB	OE=0, Rc=0
divwu.	RT, RA, RB	OE=0, Rc=1
divwuo	RT, RA, RB	OE=1, Rc=0
divwuo.	RT, RA, RB	OE=1, Rc=1

31	RT	RA	RB	OE	459	Rc
0	6	11	16	21 22		31

$$(RT) \leftarrow (RA) \div (RB)$$

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]_{LT, GT, EQ} are also undefined. The invalid division operation also sets XER[OV, SO] to 1 if the OE field contains 1.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} 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 zero.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

eieio

The **eieio** instruction ensures that all loads and stores preceding **eieio** complete with respect to main storage before any loads and stores following **eieio** access main storage.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Registers Altered

- None

Invalid Instruction Forms

- Reserved fields

Programming Note

Architecturally, **eieio** orders storage access, not instruction completion. Therefore, non-storage operations after **eieio** could complete before storage operations that were before **eieio**. The **sync** instruction guarantees ordering of both instruction completion and storage access. For the PPC405 core, the **eieio** instruction is implemented to behave as a **sync** instruction.

To write code that is portable between various PowerPC implementations, programmers should use the mnemonic that corresponds to the desired behavior.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

eqv

Equivalent

eqv RA, RS, RB Rc=0
eqv. RA, RS, RB Rc=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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

extsb RA, RS Rc=0
extsb. 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]_{LT, GT, EQ, SO} if Rc contains 1

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

extsh

Extend Sign Halfword

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]_{LT, GT, EQ, SO} if Rc contains 1


Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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 cachable in the ICCR.

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 operations use MSR[DR], not MSR[IR], to determine translation of their operands.

When data translation is disabled, cachability for the EA of the operand of instruction cache operations is determined by the ICCR, not the DCCR.

Exceptions

Instruction storage exceptions and instruction-side TLB miss exceptions are associated with instruction *fetching*, not with instruction execution. Exceptions that occur during the *execution* of instruction cache operations cause data-side exceptions (data storage exceptions and data TLB miss exceptions).

This instruction is considered a “load” with respect to data storage exceptions. See “Data Storage Interrupt” on page 5-16.

This instruction is considered a “load” with respect to data address compare (DAC) debug exceptions.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

icbt

Instruction Cache Block Touch

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 is marked as cachable, the instruction block is loaded into the instruction cache.

If the instruction block at the EA is in the instruction cache, or if the EA is marked as non-cachable, 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

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 operations use MSR[DR], not MSR[IR], to determine translation of their operands. When data translation is disabled, cachability for the effective address of the operand of instruction cache operations is determined by the ICCR, not the DCCR.

Exceptions

Instruction storage exceptions and instruction-side TLB miss exceptions are associated with instruction *fetching*, not with instruction execution. Exceptions occurring during *execution* of instruction cache operations cause data storage and data TLB miss exceptions.


If the execution of an **icbt** instruction would cause a data TLB miss exception, no operation is performed and no exception occurs.

This instruction is considered a “load” with respect to protection exceptions, but cannot cause data storage exceptions. This instruction is also considered a “load” with respect to data address compare (DAC) debug exceptions.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

iccci RA, RB



EA \leftarrow (RA|0) + (RB)
 ICCI(ICU cache array)

This instruction invalidates the entire ICU cache array. The EA is not used; previous implementations have used the EA for protection checks. The instruction form 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 cache tag array before enabling the cache. Cachability can then be enabled.


Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

icread

Instruction Cache Read

icread RA, RB



$EA \leftarrow (RA|0) + (RB)$
if $((CCR0[CIS] = 0) \wedge (CCR0[CWS] = 0))$ then $(ICDBDR) \leftarrow$ (i-cache data, way A)
if $((CCR0[CIS] = 0) \wedge (CCR0[CWS] = 1))$ then $(ICDBDR) \leftarrow$ (i-cache data, way B)
if $((CCR0[CIS] = 1) \wedge (CCR0[CWS] = 0))$ then $(ICDBDR) \leftarrow$ (i-cache tag, way A)
if $((CCR0[CIS] = 1) \wedge (CCR0[CWS] = 1))$ then $(ICDBDR) \leftarrow$ (i-cache tag, way B)

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.

This instruction is a debugging tool for reading the instruction cache entries for the congruence class specified by $EA_{18:26}$, unless no cache array is present. The cache information is read into the Instruction Cache Debug Data Register (ICDBDR), from where it can be read into a GPR using the extended mnemonic **mficbdr**.

If $CCR0[CIS] = 0$, the information is a word of instruction cache data from the addressed line. The word is specified by $EA_{27:29}$. If $CCR0[CWS] = 0$, the data is from the A-way, otherwise from the B-way.

If $(CCR0[CIS] = 1)$, the information is a cache tag from the addressed congruence class. If $(CCR0[CWS] = 0)$, the tag is from the A-way, otherwise from the B-way.

Instruction cache tag information is placed in the ICDBDR as shown:

0:21	TAG	Cache Tag
22:26		Reserved
27	V	Cache Line Valid 0 Not valid 1 Valid
28:30		Reserved
31	LRU	Least Recently Used (LRU) 0 A-way LRU 1 B-way LRU

If instruction bit 31 contains 1, the contents of $CR[CR0]$ are undefined.

Registers Altered

- ICDBDR

Invalid Instruction Forms

- Reserved fields

Programming Note

Execution of this instruction is privileged.

icread

Instruction Cache Read

The instruction pipeline does not automatically wait for data from **icread** to arrive at the ICDBDR before attempting to use the contents of the ICDBDR. Therefore, insert an **isync** instruction between **icread** and **mficdbdr**.

```
icread r5,r6 # read cache information
isync      # ensure completion of icread
mficdbdr r7 # move information to GPR
```

Instruction cache operations use MSR[DR], not MSR[IR], to determine translation of their operands. When data translation is disabled, cachability for the EA of the operand of instruction cache operations is determined by the ICCR, not the DCCR.

Exceptions

Instruction storage exceptions and instruction-side TLB miss exceptions are associated with instruction *fetching*, not with instruction execution. Exceptions that occur during the *execution* of instruction cache operations cause data-side exceptions (data storage exceptions and data TLB miss exceptions).

The execution of **icread** can cause a data TLB miss exception, at the specified EA, regardless of the non-specific intent of that EA.

This instruction is considered a “load” and cannot cause a data storage exception.

This instruction is considered a “load” with respect to data address compare (DAC) debug exceptions, but will not cause DAC debug events.


Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

isync

Instruction Synchronize

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.

No subsequent instructions are initiated by the processor until **isync** completes. Finally, execution of **isync** causes the processor to discard any prefetched instructions, with the effect that subsequent instructions are fetched and executed in the context established by the instructions preceding **isync**.

isync has no effect on caches.

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 2-33.

The following code example illustrates the necessary steps for self-modifying code. This example assumes that `addr1` is both data and instruction cachable.

```
    stw      regN, addr1      # data in regN is to become an instruction at addr1
    dcbst   addr1            # forces data from the data cache to memory
    sync    addr1            # wait until the data actually reaches the memory
    icbi    addr1            # the previous value at addr1 might already be in
                             # the instruction cache; invalidate in the cache
    isync   addr1            # the previous value at addr1 might already have been
                             # pre-fetched into the queue; invalidate the queue
                             # so that the instruction must be re-fetched
```

Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lbzu

Load Byte and Zero with Update

lbzu RT, D(RA)



$EA \leftarrow (RA|0) + EXTS(D)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow {}^{24}0 \parallel MS(EA,1)$

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 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lbzux RT, RA, RB



$EA \leftarrow (RA|0) + (RB)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow {}^{24}0 \parallel MS(EA,1)$

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 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lbzx

Load Byte and Zero Indexed

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lha RT, D(RA)
$$EA \leftarrow (RA|0) + \text{EXTS}(D)$$
$$(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$$

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhau

Load Halfword Algebraic with Update

lhau RT, D(RA)



$EA \leftarrow (RA) + EXTS(D)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow EXTS(MS(EA,2))$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhaux RT, RA, RB



$EA \leftarrow (RA) + (RB)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow EXTS(MS(EA,2))$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. 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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhax

Load Halfword Algebraic Indexed

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhbrx RT, RA, RB

$$EA \leftarrow (RA|0) + (RB)$$

$$(RT) \leftarrow {}^{16}0 \parallel MS(EA + 1, 1) \parallel MS(EA, 1)$$

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 halfword at the EA is byte-reversed. 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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhz

Load Halfword and Zero

lhz RT, D(RA)



$$EA \leftarrow (RA|0) + \text{EXTS}(D)$$
$$(RT) \leftarrow {}^{16}0 \parallel \text{MS}(EA,2)$$

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhzu RT, D(RA)



$$EA \leftarrow (RA) + \text{EXTS}(D)$$

$$(RA) \leftarrow EA$$

$$(RT) \leftarrow {}^{16}0 \parallel \text{MS}(EA,2)$$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. 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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lhzux

Load Halfword and Zero with Update Indexed

lhzux **RT, RA, RB**



$EA \leftarrow (RA) + (RB)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow {}^{16}0 \parallel MS(EA,2)$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Imw

Load Multiple Word

Imw RT, D(RA)



```
EA ← (RA|0) + EXTS(D)
r ← RT
do while r ≤ 31
  if ((r ≠ RA) ∨ (r = 31)) then
    (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). Register RA is not altered by this instruction (unless RA is GPR(31), which is an invalid form of this instruction). The word which would have been placed into register RA is discarded.

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.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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
        if ((r ≠ RA) ∨ (r = RFINAL)) then
            (GPR(r)) ← 0
    if ((r ≠ RA) ∨ (r = RFINAL)) then
        (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_{FINAL}. Register RA is not altered (unless RA = R_{FINAL}, an invalid form of this instruction). Bytes which would have been loaded into register RA are discarded.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Registers Altered

- RT and subsequent GPRs as described above.

lswi

Load String Word Immediate


Invalid Instruction Forms

- Reserved fields
- RA is in the range of registers to be loaded
- RA = RT = 0

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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
    if (((r ≠ RA) ∧ (r ≠ RB)) ∨ (r = RFINAL)) then
      (GPR(r)) ← 0
  if (((r ≠ RA) ∧ (r ≠ RB)) ∨ (r = RFINAL)) then
    (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_{FINAL}. Register RA is not altered (unless RA = R_{FINAL}, which is an invalid form of this instruction). Register RB is not altered (unless RB = R_{FINAL}, which is an invalid form of this instruction). Bytes which would have been loaded into registers RA or RB are discarded.

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.

lswx

Load String Word Indexed

Invalid Instruction Forms

- Reserved fields
- RA or RB is in the range of registers to be loaded.
- RA = RT = 0

Programming Note

If XER[TBC] = 0, the contents of register RT are unchanged and **lswx** is treated as a no-op.

The PowerPC Architecture states that, if XER[TBC] = 0 and if the EA is such that a precise data exception would normally occur (if not for the zero length), **lswx** is treated as a no-op and the precise exception will not occur. Data storage exceptions and alignment exceptions are examples of precise data exceptions.


However, the PowerPC Architecture makes no statement regarding imprecise exceptions related to **lswx** with XER[TBC] = 0. The PPC405 core generates an imprecise exception (machine check) on this instruction when all of the following conditions are true:

- The instruction passes all protection bounds checking
- The address is cachable
- The address is passed to the data cache
- The address misses in the data cache (resulting in a line fill request)
- The address encounters some form of bus error

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lwarx RT, RA, RB



```
EA ← (RA|0) + (RB)
RESERVE ← 1
(RT) ← 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

lwarx and the **stwcx.** instruction should be 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 asynchronous processes. Only **lwarx** can set the reservation bit to 1. **stwcx.** sets the reservation bit to 0 upon its completion, whether or not **stwcx.** sent (RS) to memory. CR[CR0]_{EQ} 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 register as required
      stwcx. # attempt to store semaphore; reset reservation
      bne loop # an asynchronous process has intervened; try again
```

If the asynchronous process in the code example had paired **lwarx** with a store other than **stwcx.**, the reservation bit would not have been cleared in the asynchronous process, and the code example would have overwritten the semaphore.

Exceptions

An alignment exception occurs if the EA is not word-aligned.


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lwbrx

Load Word Byte-Reverse Indexed

lwbrx RT, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$
$$(RT) \leftarrow MS(EA+3,1) \parallel MS(EA+2,1) \parallel MS(EA+1,1) \parallel MS(EA,1)$$

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 word at the EA is byte-reversed: the least significant byte becomes the most significant byte, the next least significant byte becomes the next most significant byte, and so on. The resulting word 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lwzu

Load Word and Zero with Update

lwzu RT, D(RA)



$EA \leftarrow (RA) + \text{EXTS}(D)$
 $(RA) \leftarrow EA$
 $(RT) \leftarrow \text{MS}(EA, 4)$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lwzux RT, RA, RB

$$EA \leftarrow (RA) + (RB)$$

$$(RA) \leftarrow EA$$

$$(RT) \leftarrow MS(EA,4)$$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. 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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

lwzx

Load Word and Zero Indexed

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

macchw	RT, RA, RB	OE=0, Rc=0
macchw.	RT, RA, RB	OE=0, Rc=1
macchwo	RT, RA, RB	OE=1, Rc=0
macchwo.	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ 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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

macchws

Multiply Accumulate Cross Halfword to Word Saturate Signed

macchws	RT, RA, RB	OE=0, Rc=0
macchws.	RT, RA, RB	OE=0, Rc=1
macchwso	RT, RA, RB	OE=1, Rc=0
macchwso.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	236	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ signed
 $temp_{0:32} \leftarrow prod_{0:31} + (RT)$
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 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 the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

macchwsu	RT, RA, RB	OE=0, Rc=0
macchwsu.	RT, RA, RB	OE=0, Rc=1
macchwsuo	RT, RA, RB	OE=1, Rc=0
macchwsuo.	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 and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is greater than $2^{32} - 1$, the value stored in RT is $2^{32} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

macchwu

Multiply Accumulate Cross Halfword to Word Modulo Unsigned

macchwu	RT, RA, RB	OE=0, Rc=0
macchwu.	RT, RA, RB	OE=0, Rc=1
macchwuo	RT, RA, RB	OE=1, Rc=0
macchwuo.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	140	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$

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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

machhw	RT, RA, RB	OE=0, Rc=0
machhw.	RT, RA, RB	OE=0, Rc=1
machhwo	RT, RA, RB	OE=1, Rc=0
machhwo.	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ 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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

machhws

Multiply Accumulate High Halfword to Word Saturate Signed

machhws	RT, RA, RB	OE=0, Rc=0
machhws.	RT, RA, RB	OE=0, Rc=1
machhwsO	RT, RA, RB	OE=1, Rc=0
machhwsO.	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 and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

machhwsu

Multiply Accumulate High Halfword to Word Saturate Unsigned

machhwsu	RT, RA, RB	OE=0, Rc=0
machhwsu.	RT, RA, RB	OE=0, Rc=1
machhwsuo	RT, RA, RB	OE=1, Rc=0
machhwsuo.	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow (temp_{1:32} \vee^{32} temp_0)$

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 and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is greater than $2^{32} - 1$, the value stored in RT is $2^{32} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1


Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

machhwu

Multiply Accumulate High Halfword to Word Modulo Unsigned

machhwu	RT, RA, RB	OE=0, Rc=0
machhwu.	RT, RA, RB	OE=0, Rc=1
machhwuo	RT, RA, RB	OE=1, Rc=0
machhwuo.	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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.


Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

maclhw	RT, RA, RB	OE=0, Rc=0
maclhw.	RT, RA, RB	OE=0, Rc=1
maclhwo	RT, RA, RB	OE=1, Rc=0
maclhwo.	RT, RA, RB	OE=1, Rc=1



$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ 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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

maclhws

Multiply Accumulate Low Halfword to Word Saturate Signed

maclhws	RT, RA, RB	OE=0, Rc=0
maclhws.	RT, RA, RB	OE=0, Rc=1
maclhws0	RT, RA, RB	OE=1, Rc=0
maclhws0.	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 and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

maclhwsu	RT, RA, RB	OE=0, Rc=0
maclhwsu.	RT, RA, RB	OE=0, Rc=1
maclhwsuo	RT, RA, RB	OE=1, Rc=0
maclhwsuo.	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 and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is greater than $2^{32} - 1$, the value stored in RT is $2^{32} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mac|hwu

Multiply Accumulate Low Halfword to Word Modulo Unsigned

mac hwu	RT, RA, RB	OE=0, Rc=0
mac hwu.	RT, RA, RB	OE=0, Rc=1
mac hwuo	RT, RA, RB	OE=1, Rc=0
mac hwuo.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	396	Rc
0	6	11	16	21 22		31

$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned

$temp_{0:32} \leftarrow prod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$


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 the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mcrxr

Move to Condition Register from XER

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.

This transfer is positional, by bit number, so the mnemonics associated with each bit are changed. See Table 9-18 for clarification.

Table 9-18. Transfer Bit Mnemonic Assignment

Bit	XER Usage	CR Usage
0	SO	LT
1	OV	GT
2	CA	EQ
3	Reserved	SO

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mfdcr

Move from Device Control Register

mfdcr RT, DCRN



$DCRN \leftarrow DCRF_{5:9} \parallel DCRF_{0:4}$
 $(RT) \leftarrow (DCR(DCRN))$

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
- Invalid DCRF values


Programming Note

Execution of this instruction is privileged.

The DCR number (DCRN) specified in the assembler language coding of **mfdcr** refers to a DCR number. The assembler handles the unusual register number encoding to generate the DCRF field.

Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

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.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

mfspr

Move From Special Purpose Register

mfspr RT, SPRN



$SPRN \leftarrow SPRF_{5:9} \parallel SPRF_{0:4}$
 $(RT) \leftarrow (SPR(SPRN))$

The contents of the SPR specified by the SPRF field are placed into register RT. See “Special Purpose Registers” on page 10-2 for a listing of SPR mnemonics and corresponding SPRN and SPRF 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 Mode Operation” on page 2-30.

The SPR number (SPRN) specified in the assembler language coding of **mfspr** refers to an SPR number (see “Special Purpose Registers” on page 10-2 for a list of SPRN values). The assembler handles the unusual register number encoding to generate the SPRF field. Also, see “Privileged SPRs” on page 2-32 for information about privileged SPRs.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-19. Extended Mnemonics for mf spr

Mnemonic	Operands	Function	Other Registers Changed
mfccr0 mfctr mfdac1 mfdac2 mfdear mfdbcr0 mfdbcr1 mfdbsr mfdccr mfdcwr mfdvc1 mfdvc2 mfesr mfevpr mfiac1 mfiac2 mfiac3 mfiac4 mficcr mfidbdr mflr mfpid mfpit mfpvr mfsgr mfsler mfsprg0 mfsprg1 mfsprg2 mfsprg3 mfsprg4 mfsprg5 mfsprg6 mfsprg7 mfsrr0 mfsrr1 mfsrr2 mfsrr3 mfsu0r mftcr mftsr mfxer mfzpr	RT	Move from special purpose register SPRN. <i>Extended mnemonic for</i> mf spr RT,SPRN See “Special Purpose Registers” on page 10-2 for a list of valid SPRN values.	

mftb

Move From Time Base

mftb RT, TBRN



$$\begin{aligned} \text{TBRN} &\leftarrow \text{TBRF}_{5:9} \parallel \text{TBRF}_{0:4} \\ (\text{RT}) &\leftarrow (\text{TBR}(\text{TBRN})) \end{aligned}$$

The contents of the time base register (TBR) specified by the TBRF field are placed into register RT. The following table lists the TBRN and TBRF values.

Table 9-20. Extended Mnemonics for mftb

Register Mnemonic	Register Name	TBRN		TBRF	Access
		Decimal	Hex		
TBL	Time Base Lower	268	0x10C	0x188	Read-only
TBU	Time Base Upper	269	0x10D	0x1A8	Read-only

If TBRN is a value other than those listed in the table, the results are boundedly undefined.

Registers Altered

- RT

Invalid Instruction Forms

- Reserved fields
- Invalid TBRF values

Programming Notes

The mnemonic **mftb** serves as both a hardware mnemonic and an extended mnemonic. The assembler recognizes an **mftb** mnemonic having two operands as the hardware form; an **mftb** mnemonic having one operand is recognized as the extended form.

The TBR number (TBRN) specified in the assembler language coding of the **mftb** instruction refers to a TBR number listed in the preceding table. The assembler handles the unusual register number encoding to generate the TBRF field.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Virtual Environment.


Table 9-21. Extended Mnemonics for mftb

Mnemonic	Operands	Function	Other Registers Altered
mftb	RT	Move the contents of TBL into RT. <i>Extended mnemonic for mftb RT,TBL</i>	
mftbu	RT	Move the contents of TBU into RT. <i>Extended mnemonic for mftb RT,TBU</i>	

mtcrf

Move to Condition Register Fields

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:

FXM Bit Number	Bits Controlled
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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-22. Extended Mnemonics for mtcrrf

Mnemonic	Operands	Function	Other Registers Altered
mtcr	RS	Move to CR. <i>Extended mnemonic for mtcrrf 0xFF,RS</i>	

mtdcr DCRN, RS

$$\text{DCRN} \leftarrow \text{DCRF}_{5:9} \parallel \text{DCRF}_{0:4}$$

$$(\text{DCR}(\text{DCRN})) \leftarrow (\text{RS})$$

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
- Invalid DCRF values

Programming Note

Execution of this instruction is privileged.

The DCR number (DCRN) specified in the assembler language coding of **mtdcr** refers to a DCR number. The assembler handles the unusual register number encoding to generate the DCRF field.

Architecture Note

This instruction is implementation-specific and may not be portable to other implementations.

mtmsr

Move To Machine State Register

mtmsr RS



$(MSR) \leftarrow (RS)$

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.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

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” on page 10-2 for a listing of SPR mnemonics and corresponding SPRN and SPRF 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 is a 1. See “Privileged SPRs” on page 2-32 for more information.

The SPR number (SPRN) specified in the assembler language coding of the **mtspr** instruction refers to an SPR number (see “Special Purpose Registers” on page 10-2 for a list of SPRN values). The assembler handles the unusual register number encoding to generate the SPRF field.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mtspr

Move To Special Purpose Register

Table 9-23. Extended Mnemonics for mtspr

Mnemonic	Operands	Function	Other Registers Altered
mtccr0 mtctr mtdac1 mtdac2 mtdbcr0 mtdbcr1 mtdbsr mtdccr mtdcwr mtdear mtdvc1 mtdvc2 mtesr mtevpr mtiac1 mtiac2 mtiac3 mtiac4 mticcr mticdbdr mtr mtpid mtpit mtpvr mtsgr mtsler mtsprg0 mtsprg1 mtsprg2 mtsprg3 mtsprg4 mtsprg5 mtsprg6 mtsprg7 mtsrr0 mtsrr1 mtsrr2 mtsrr3 mtsuo0r mttcr mttsr mtxer mtzpr	RS	Move to special purpose register SPRN. <i>Extended mnemonic for</i> mtspr SPRN,RS See “Special Purpose Registers” on page 10-2 for a list of valid SPRN values.	

mulchw RT, RA, RB Rc=0
mulchw. RT, RA, RB Rc=1

4	RT	RA	RB	168	Rc
0	6	11	16	21	31

$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ signed

The low-order halfword of RA is multiplied by the high-order halfword of RB. The resulting signed product replaces the contents of RT.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mulchwu

Multiply Cross Halfword to Word Unsigned

mulchwu RT, RA, RB Rc=0
mulchwu. RT, RA, RB Rc=1

4	RT	RA	RB	136	Rc
0	6	11	16	21	31

$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15} \text{ unsigned}$$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The resulting unsigned product replaces the contents of RT.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mulhww RT, RA, RB Rc=0
mulhww. RT, RA, RB Rc=1

4	RT	RA	RB	40	Rc
0	6	11	16	21	31

$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ signed

The high-order halfword of RA is multiplied by the high-order halfword of RB. The resulting signed product replaces the contents of RT.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1


Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mulhhwu

Multiply High Halfword to Word Unsigned

mulhhwu RT, RA, RB Rc=0
mulhhwu. 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. The resulting unsigned product replaces the contents of RT.


Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mulhw RT, RA, RB Rc=0
mulhw. 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]_{LT, GT, EQ, SO} 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.


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mulhwu

Multiply High Word Unsigned

mulhwu RT, RA, RB Rc=0
mulhwu. RT, RA, RB Rc=1



$prod_{0:63} \leftarrow (RA) \times (RB) \text{ unsigned}$
 $(RT) \leftarrow prod_{0:31}$

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]_{LT, GT, EQ, SO} 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.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mullhw RT, RA, RB Rc=0
mullhw. 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. The resulting signed product replaces the contents of RT.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mullhwu

Multiply Low Halfword to Word Unsigned

mullhwu RT, RA, RB OE=0, Rc=0
mullhwu. RT, RA, RB OE=0, Rc=1

4	RT	RA	RB	392	Rc
0	6	11	16	21	31

$$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31} \text{ unsigned}$$

The low-order halfword of RA is multiplied by the low-order halfword of RB. The resulting unsigned product replaces the contents of RT.


Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

mulli RT, RA, IM



$$\text{prod}_{0:47} \leftarrow (\text{RA}) \times \text{EXTS}(\text{IM}) \text{ signed}$$

$$(\text{RT}) \leftarrow \text{prod}_{16:47}$$

The 48-bit product of register RA and the sign-extended IM field is formed. Both register RA and the IM field are interpreted as signed quantities. 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.


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

mullw

Multiply Low Word

mullw	RT, RA, RB	OE=0, Rc=0
mullw.	RT, RA, RB	OE=0, Rc=1
mullwo	RT, RA, RB	OE=1, Rc=0
mullwo.	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]_{LT, GT, EQ, SO} 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 is correct only if the operands are regarded as signed numbers.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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]_{LT, GT, EQ, SO} if Rc contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

neg

Negate

neg	RT, RA	OE=0, Rc=0
neg.	RT, RA	OE=0, Rc=1
nego	RT, RA	OE=1, Rc=0
nego.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE=1

Invalid Instruction Forms

- Reserved fields

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

nmacchw	RT, RA, RB	OE=0, Rc=0
nmacchw.	RT, RA, RB	OE=0, Rc=1
nmacchwo	RT, RA, RB	OE=1, Rc=0
nmacchwo.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	174	Rc
0	6	11	16	21 22		31

$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15})$ signed

$temp_{0:32} \leftarrow nprod_{0:31} + (RT)$

$(RT) \leftarrow temp_{1:32}$

The low-order halfword of RA is multiplied by the high-order halfword of RB. The negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nmacchws

Negative Multiply Accumulate Cross Halfword to Word Saturate Signed

nmacchws	RT, RA, RB	OE=0, Rc=0
nmacchws.	RT, RA, RB	OE=0, Rc=1
nmacchwso	RT, RA, RB	OE=1, Rc=0
nmacchwso.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	238	Rc
0	6	11	16	21 22		31

$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15} \text{ signed})$
 $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$
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 low-order halfword of RA is multiplied by the high-order halfword of RB. The negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nmachhw	RT, RA, RB	OE=0, Rc=0
nmachhw.	RT, RA, RB	OE=0, Rc=1
nmachhwo	RT, RA, RB	OE=1, Rc=0
nmachhwo.	RT, RA, RB	OE=1, Rc=1



$$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15}) \text{ signed}$$

$$temp_{0:32} \leftarrow nprod_{0:31} + (RT)$$

$$(RT) \leftarrow temp_{1:32}$$

The high-order halfword of RA is multiplied by the high-order halfword of RB. The negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1


Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nmachhws

Negative Multiply Accumulate High Halfword to Word Saturate Signed

nmachhws	RT, RA, RB	OE=0, Rc=0
nmachhws.	RT, RA, RB	OE=0, Rc=1
nmachhwsO	RT, RA, RB	OE=1, Rc=0
nmachhwsO.	RT, RA, RB	OE=1, Rc=1



$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ signed
 $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$
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 negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register.

If a result does not overflow (i.e., it is accurately representable in 32 bits), the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.


Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nmac1hw	RT, RA, RB	OE=0, Rc=0
nmac1hw.	RT, RA, RB	OE=0, Rc=1
nmac1hwo	RT, RA, RB	OE=1, Rc=0
nmachlwo.	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 negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register. The contents of RT are replaced by the low-order 32 bits of the temporary register.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nmaclhws

Negative Multiply Accumulate High Halfword to Word Saturate Signed

nmaclhws	RT, RA, RB	OE=0, Rc=0
nmaclhws.	RT, RA, RB	OE=0, Rc=1
nmaclhws0	RT, RA, RB	OE=1, Rc=0
nmaclhws0.	RT, RA, RB	OE=1, Rc=1

4	RT	RA	RB	OE	494	Rc
0	6	11	16	21 22		31

$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{16:31})$ 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 low-order halfword of RA is multiplied by the low-order halfword of RB. The negated signed product is summed with the contents of RT and the sum is stored in a 33-bit temporary register.

If a result does not overflow, the low-order 32 bits of the temporary register are stored in RT.

If a result overflows, the returned result is the nearest representable value. Thus, if a result is less than -2^{31} , the value stored in RT is -2^{31} . Likewise, if a result is greater than $2^{31} - 1$, the value stored in RT is $2^{31} - 1$.

Registers Altered

- RT
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the Multiply-Accumulate instruction set extensions and complies with the architectural requirements for APUs of the IBM PowerPC Embedded Environment. As such, it is not part of the PowerPC Architecture, nor is it part of the IBM PowerPC Embedded Environment. Programs that use this instruction may not be portable to other implementations.

nor RA, RS, RB Rc=0
nor. RA, RS, RB Rc=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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-24. Extended Mnemonics for nor, nor.

Mnemonic	Operands	Function	Other Registers Altered
not	RA, RS	Complement register. (RA) ← ¬(RS) <i>Extended mnemonic for nor RA,RS,RS</i>	
not.		<i>Extended mnemonic for nor. RA,RS,RS</i>	CR[CR0]

or

OR

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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-25. Extended Mnemonics for or, or.

Mnemonic	Operands	Function	Other Registers Altered
mr	RT, RS	Move register. (RT) ← (RS) <i>Extended mnemonic for or RT,RS,RS</i>	
mr.		<i>Extended mnemonic for or. RT,RS,RS</i>	CR[CR0]

orc

OR with Complement

orc RA, RS, RB Rc=0
orc. RA, RS, RB Rc=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]_{LT, GT, EQ, SO} if Rc contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

ori

OR Immediate

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-26. Extended Mnemonics for ori

Mnemonic	Operands	Function	Other Registers Changed
nop		Preferred no-op; triggers optimizations based on no-ops. <i>Extended mnemonic for</i> ori 0,0,0	

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

rfci

Return From Critical Interrupt

rfci



(PC) ← (SRR2)
(MSR) ← (SRR3)

The program counter (PC) is restored with the contents of SRR2 and the MSR is restored with the contents of SRR3.

Instruction execution returns to the address contained in the PC.

Registers Altered


- MSR

Programming Note

Execution of this instruction is privileged and context-synchronizing.

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

rfi

(PC) ← (SRR0)
(MSR) ← (SRR1)

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.

Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

rlwimi

Rotate Left Word Immediate then Mask Insert

rlwimi RA, RS, SH, MB, ME Rc=0
rlwimi. RA, RS, SH, MB, ME Rc=1

20	RS	RA	SH	MB	ME	Rc
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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-27. Extended Mnemonics for rlwimi, rlwimi.

Mnemonic	Operands	Function	Other Registers Altered
inslwi	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> rlwimi RA,RS,32-b,b,b+n-1	
inslwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b,b,b+n-1	CR[CR0]
insrwi	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> rlwimi RA,RS,32-b-n,b,b+n-1	
insrwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b-n,b,b+n-1	CR[CR0]

rlwinm RA, RS, SH, MB, ME Rc=0
rlwinm. RA, RS, SH, MB, ME Rc=1

21	RS	RA	SH	MB	ME	Rc
0	6	11	16	21	26	31

$r \leftarrow \text{ROTL}((RS), SH)$
 $m \leftarrow \text{MASK}(MB, ME)$
 $(RA) \leftarrow r \wedge m$

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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-28. Extended Mnemonics for rlwinm, rlwinm.

Mnemonic	Operands	Function	Other Registers Altered
clrlwi	RA, RS, n	Clear left immediate. ($n < 32$) $(RA)_{0:n-1} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,n,31	
clrlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,n,31	CR[CR0]
clrlslwi	RA, RS, b, n	Clear left and shift left immediate. $(n \leq b < 32)$ $(RA)_{b-n:31-n} \leftarrow (RS)_{b:31}$ $(RA)_{32-n:31} \leftarrow {}^n0$ $(RA)_{0:b-n-1} \leftarrow {}^{b-n}0$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,b-n,31-n	
clrlslwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,b-n,31-n	CR[CR0]

rlwinm

Rotate Left Word Immediate then AND with Mask

Table 9-28. Extended Mnemonics for rlwinm, rlwinm. (continued)

Mnemonic	Operands	Function	Other Registers Altered
clrrwi	RA, RS, n	Clear right immediate. ($n < 32$) $(RA)_{32-n:31} \leftarrow ^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,0,31-n	
clrrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,0,31-n	CR[CR0]
extlwi	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> rlwinm RA,RS,b,0,n-1	
extlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,b,0,n-1	CR[CR0]
extrwi	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> rlwinm RA,RS,b+n,32-n,31	
extrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,b+n,32-n,31	CR[CR0]
rotlwi	RA, RS, n	Rotate left immediate. $(RA) \leftarrow \text{ROTL}((RS), n)$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,0,31	
rotlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31	CR[CR0]
rotrwi	RA, RS, n	Rotate right immediate. $(RA) \leftarrow \text{ROTL}((RS), 32-n)$ <i>Extended mnemonic for</i> rlwinm RA,RS,32-n,0,31	
rotrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,0,31	CR[CR0]
slwi	RA, RS, n	Shift left immediate. ($n < 32$) $(RA)_{0:31-n} \leftarrow (RS)_{n:31}$ $(RA)_{32-n:31} \leftarrow ^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,0,31-n	
slwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31-n	CR[CR0]

Table 9-28. Extended Mnemonics for rlwinm, rlwinm. (continued)

Mnemonic	Operands	Function	Other Registers Altered
srwi	RA, RS, n	Shift right immediate. ($n < 32$) $(RA)_{n:31} \leftarrow (RS)_{0:31-n}$ $(RA)_{0:n-1} \leftarrow ^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,32-n,n,31	
srwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,n,31	CR[CR0]

rlwnm

Rotate Left Word then AND with Mask

rlwnm RA, RS, RB, MB, ME Rc=0
rlwnm. 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_{27:31}. 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]_{LT, GT, EQ, SO} if Rc contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-29. Extended Mnemonics for rlwnm, rlwnm.

Mnemonic	Operands	Function	Other Registers Altered
rotlw	RA, RS, RB	Rotate left. $(RA) \leftarrow \text{ROTL}((RS), (RB)_{27:31})$ <i>Extended mnemonic for</i> rlwnm RA,RS,RB,0,31	
rotlw.		<i>Extended mnemonic for</i> rlwnm. RA,RS,RB,0,31	CR[CR0]

sc



$(SRR1) \leftarrow (MSR)$
 $(SRR0) \leftarrow (PC)$
 $PC \leftarrow EVPR_{0:15} \parallel 0x0C00$
 $(MSR[WE, EE, PR, DR, IR]) \leftarrow 0$

A system call exception is generated. 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 exception vector address. The exception vector address is calculated by concatenating the high halfword of the Exception Vector Prefix Register (EVPR) to the left of 0x0C00.

The MSR[WE, EE, PR, DR, IR] bits are set to 0.

Program execution continues at the new address in the PC.

The **sc** instruction is context synchronizing.

Registers Altered

- SRR0
- SRR1
- MSR[WE, EE, PR, DR, IR]

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

slw

Shift Left Word

slw RA, RS, RB Rc=0
slw. RA, RS, RB Rc=1



```
n ← (RB)27:31
r ← ROTL((RS), n)
if (RB)26 = 0 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_{27:31}. 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.

If RB₂₆ = 1, register RA is set to zero.

Registers Altered

- RA
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sraw RA, RS, RB Rc=0
sraw. RA, RS, RB Rc=1

31	RS	RA	RB	792	Rc
0	6	11	16	21	31

```

n ← (RB)27:31
r ← ROTL((RS), 32 - n)
if (RB)26 = 0 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_{27:31}. Bits shifted out of the least significant bit are lost. Register RS₀ 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.

If bit 26 of register RB contains 1, register RA and XER[CA] are set to bit 0 of register RS.

Registers Altered

- RA
- XER[CA]
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

srawi

Shift Right Algebraic Word Immediate

srawi RA, RS, SH Rc=0
srawi. 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 RS₀ 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]_{LT, GT, EQ, SO} if Rc contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

srw

Shift Right Word

srw RA, RS, RB Rc=0
srw. RA, RS, RB Rc=1



```

n ← (RB)27:31
r ← ROTL((RS), 32 – n)
if (RB)26 = 0 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_{27:31}. 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.

If bit 26 of register RB contains a one, register RA is set to 0.

Registers Altered

- RA
- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stb

Store Byte

stb RS, D(RA)



$EA \leftarrow (RA|0) + \text{EXTS}(D)$

$MS(EA, 1) \leftarrow (RS)_{24:31}$

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stbu RS, D(RA)

$$EA \leftarrow (RA) + \text{EXTS}(D)$$

$$MS(EA, 1) \leftarrow (RS)_{24:31}$$

$$(RA) \leftarrow EA$$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The EA is placed into register RA.

The least significant byte of register RS is stored into the byte at the EA.

Registers Altered

- RA

Invalid Instruction Forms

RA = 0

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stbux

Store Byte with Update Indexed

stbux RS, RA, RB



$EA \leftarrow (RA) + (RB)$
 $MS(EA, 1) \leftarrow (RS)_{24:31}$
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. The EA is placed into register RA.

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


- RA

Invalid Instruction Forms

- Reserved fields
- RA = 0

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sth

Store Halfword

sth RS, D(RA)



$$EA \leftarrow (RA|0) + \text{EXTS}(D)$$

$$MS(EA, 2) \leftarrow (RS)_{16:31}$$

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sthbrx RS, RA, RB

$$EA \leftarrow (RA|0) + (RB)$$

$$MS(EA, 2) \leftarrow (RS)_{24:31} \parallel (RS)_{16:23}$$

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. The result 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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sth

Store Halfword with Update

sth RS, D(RA)



$EA \leftarrow (RA) + EXTS(D)$
 $MS(EA, 2) \leftarrow (RS)_{16:31}$
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The EA is placed into register RA.

The least significant halfword of register RS is stored into the halfword at the EA.

Registers Altered

- RA


Invalid Instruction Forms

- RA = 0

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sthux RS, RA, RB



$$EA \leftarrow (RA) + (RB)$$

$$MS(EA, 2) \leftarrow (RS)_{16:31}$$

$$(RA) \leftarrow EA$$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. The EA is placed into register RA.

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

- RA

Invalid Instruction Forms

- Reserved fields
- RA = 0


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sthx

Store Halfword Indexed

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stswi

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), wrapping to GPR(0), 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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), wrapping to GPR(0), 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

If XER[TBC] = 0, **stswx** is treated as a no-op.

The PowerPC Architecture states that if XER[TBC] = 0 and if the EA is such that a precise data exception would normally occur (if not for the zero length), **stswx** is treated as a no-op and the precise exception will not occur. Data storage exceptions and alignment exceptions are examples of precise data exceptions.

stswx

Store String Word Indexed


However, the architecture makes no statement regarding imprecise exceptions related to **stswx** when $XER[TBC] = 0$. IBM PowerPC processors generate an imprecise exception (machine check) on this instruction when all of the following conditions are true:

- The instruction passes all protection bounds checking
- The address is cachable
- The address is passed to the data cache
- The address misses in the data cache (resulting in a line fill request)
- The address encounters some form of bus error (non-configured, for example)

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stwbrx

Store Word Byte-Reverse Indexed

stwbrx RS, RA, RB



$$EA \leftarrow (RA|0) + (RB)$$

$$MS(EA, 4) \leftarrow (RS)_{24:31} \parallel (RS)_{16:23} \parallel (RS)_{8:15} \parallel (RS)_{0:7}$$

An 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 byte-reversed: the least significant byte becomes the most significant byte, the next least significant byte becomes the next most significant byte, and so on. The result is 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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stwcx. RS, RA, RB

```

EA ← (RA|0) + (RB)
if RESERVE = 1 then
    MS(EA, 4) ← (RS)
    RESERVE ← 0
    (CR[CR0]) ← 20 || 1 || XERSO
else
    (CR[CR0]) ← 20 || 0 || XERSO

```

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]_{LT, GT} are cleared
- CR[CR0]_{EQ} is set to the state of the reservation bit at the start of the instruction
- CR[CR0]_{SO} is set to the contents of the XER[SO] bit

Registers Altered

- CR[CR0]_{LT, GT, EQ, SO}

Programming Note

lwarx and the **stwcx.** instruction should be 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 asynchronous processes. Only **lwarx** can set the reservation bit to 1. **stwcx.** sets the reservation bit to 0 upon its completion, whether or not **stwcx.** sent (RS) to memory. CR[CR0]_{EQ} 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 register as required
      stwcx. # attempt to store semaphore; reset reservation
      bne loop # an asynchronous process has intervened; try again

```

If the asynchronous process in the code example had paired **lwarx** with a store other than **stwcx.**, the reservation bit would not have been cleared in the asynchronous process, and the code example would have overwritten the semaphore.

Exceptions


An alignment exception occurs if the EA is not word-aligned.

stwcx.

Store Word Conditional Indexed

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stwu RS, D(RA)
$$\begin{aligned}EA &\leftarrow (RA) + \text{EXTS}(D) \\MS(EA, 4) &\leftarrow (RS) \\(RA) &\leftarrow EA\end{aligned}$$

An effective address (EA) is formed by adding a displacement to the base address in register RA. The displacement is obtained by sign-extending the 16-bit D field to 32 bits. The EA is placed into register RA.

The contents of register RS are stored into the word at the EA.

Registers Altered

- RA

Invalid Instruction Forms

- RA = 0


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

stwux

Store Word with Update Indexed

stwux RS, RA, RB



$EA \leftarrow (RA) + (RB)$
 $MS(EA, 4) \leftarrow (RS)$
 $(RA) \leftarrow EA$

An effective address (EA) is formed by adding an index to the base address in register RA. The index is the contents of register RB. The EA is placed into register RA.

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


- RA

Invalid Instruction Forms

- Reserved fields
- RA = 0

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

subf

Subtract From

subf	RT, RA, RB	OE=0, Rc=0
subf.	RT, RA, RB	OE=0, Rc=1
subfo	RT, RA, RB	OE=1, Rc=0
subfo.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-30. Extended Mnemonics for subf, subf., subfo, subfo.

Mnemonic	Operands	Function	Other Registers Altered
sub	RT, RA, RB	Subtract (RB) from (RA). $(RT) \leftarrow \neg(RB) + (RA) + 1.$ <i>Extended mnemonic for subf RT,RB,RA</i>	
sub.		<i>Extended mnemonic for subf. RT,RB,RA</i>	CR[CR0]
subo		<i>Extended mnemonic for subfo RT,RB,RA</i>	XER[SO, OV]
subo.		<i>Extended mnemonic for subfo. RT,RB,RA</i>	CR[CR0] XER[SO, OV]

subfc

Subtract From Carrying

subfc	RT, RA, RB	OE=0, Rc=0
subfc.	RT, RA, RB	OE=0, Rc=1
subfco	RT, RA, RB	OE=1, Rc=0
subfco.	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  $\geq$  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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-31. Extended Mnemonics for subfc, subfc., subfco, subfco.

Mnemonic	Operands	Function	Other Registers Altered
subc	RT, RA, RB	Subtract (RB) from (RA). (RT) ← ¬(RB) + (RA) + 1. Place carry-out in XER[CA]. <i>Extended mnemonic for subfc RT,RB,RA</i>	
subc.		<i>Extended mnemonic for subfc. RT,RB,RA</i>	CR[CR0]
subfco		<i>Extended mnemonic for subfco RT,RB,RA</i>	XER[SO, OV]
subfco.		<i>Extended mnemonic for subfco. RT,RB,RA</i>	CR[CR0] XER[SO, OV]

subfe

Subtract From Extended

subfe	RT, RA, RB	OE=0, Rc=0
subfe.	RT, RA, RB	OE=0, Rc=1
subfeo	RT, RA, RB	OE=1, Rc=0
subfeo.	RT, RA, RB	OE=1, Rc=1

31	RT	RA	RB	OE	136	Rc
0	6	11	16	21 22		31

```
(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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

subfic

Subtract From Immediate Carrying

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]


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

subfme

Subtract from Minus One Extended

subfme	RT, RA	OE=0, Rc=0
subfme.	RT, RA	OE=0, Rc=1
subfmeo	RT, RA	OE=1, Rc=0
subfmeo.	RT, RA	OE=1, Rc=1



```
(RT) ← ¬(RA) – 1 + XER[CA]
if ¬(RA) + 0xFFFF FFFF + XER[CA] u > 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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1
- XER[CA]

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

subfze

Subtract from Zero Extended

subfze	RT, RA	OE=0, Rc=0
subfze.	RT, RA	OE=0, Rc=1
subfzeo	RT, RA	OE=1, Rc=0
subfzeo.	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]_{LT, GT, EQ, SO} if Rc contains 1
- XER[SO, OV] if OE contains 1

Invalid Instruction Forms

- Reserved fields


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

sync

Synchronize

sync



The **sync** instruction guarantees that all instructions initiated by the processor preceding **sync** will complete before **sync** completes, and that no subsequent instructions will be initiated by the processor until after **sync** completes. When **sync** completes, all storage accesses that were initiated by the processor before the **sync** instruction will have been completed with respect to all mechanisms that access storage.

If instruction bit 31 contains 1, the contents of CR[CR0] are undefined.

Registers Altered

- None.

Invalid Instruction Forms

- Reserved fields


Programming Note

Architecturally, the **eieio** instruction orders storage access, not instruction completion. Therefore, non-storage operations that follow **eieio** could complete before storage operations that precede **eieio**. The **sync** instruction guarantees ordering of instruction completion and storage access. For the PPC405 core, the **eieio** instruction is implemented to behave as a **sync** instruction.

To write code that is portable between various PowerPC implementations, programmers should use the mnemonic that corresponds to the desired behavior.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

tlbia

All of the entries in the TLB are invalidated and become unavailable for translation by clearing the valid (V) bit in the TLBHI portion of each TLB entry. The rest of the fields in the TLB entries are unmodified.

Registers Altered

- None.

Invalid Instruction Forms

- None.

Programming Note

This instruction is privileged. Translation is not required to be active during the execution of this instruction. The effects of the invalidation are not guaranteed to be visible to the programming model until the completion of a context synchronizing operation.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

tlbre

TLB Read Entry

tlbre RT, RA, WS



```
if WS4 = 1
    (RT) ← TLBLO[(RA26:31)]
else
    (RT) ← TLBHI[(RA26:31)]
    (PID) ← TID from TLB[(RA26:31)]
```

The contents of the selected TLB entry is placed into register RT (and possibly into PID).

Bits 26:31 of the contents of RA is used as an index into the TLB. If this index specifies a TLB entry that does not exist, the results are undefined.

The WS field specifies which portion (TLBHI or TLBLO) of the entry is loaded into RT. If TLBHI is being accessed, the PID SPR is set to the value of the TID field in the TLB entry.

If the WS field is not 0 or 1, 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

- RT
- PID (if WS = 0)

Invalid Instruction Forms

- Reserved fields
- Invalid WS value

Programming Notes

This instruction is privileged. Translation is not required to be active during the execution of this instruction.

The contents of RT after the execution of this instruction are interpreted as follows:

```
If WS = 0 (TLBHI):
    RT[0:21] ← EPN[0:21]
    RT[22:24] ← SIZE[0:2]
    RT[25] ← V
    RT[26] ← E
    RT[27] ← U0
    RT[28:31] ← 0
    PID[24:31] ← TID[0:7]; (note that the TID is copied to the PID, not to RT)
If WS = 1 (TLBLO):
    RT[0:21] ← RPN[0:21]
    RT[22:23] ← EX,WR
    RT[24:27] ← ZSEL[0:3]
    RT[28:31] ← WIMG
```

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.


Table 9-32. Extended Mnemonics for tlbre

Mnemonic	Operands	Function	Other Registers Altered
tlbrehi	RT, RA	Load TLBHI portion of the selected TLB entry into RT. Load the PID register with the contents of the TID field of the selected TLB entry. $(RT) \leftarrow TLBHI[(RA)]$ $(PID) \leftarrow TLB[(RA)]_{TID}$ <i>Extended mnemonic for</i> tlbre RT,RA,0	
tlbrelo	RT, RA	Load TLBLO portion of the selected TLB entry into RT. $(RT) \leftarrow TLBLO[(RA)]$ <i>Extended mnemonic for</i> tlbre RT,RA,1	

tlbsx

TLB Search Indexed

tlbsx RT, RA, RB Rc=0
tlbsx. RT, RA, RB Rc=1



```
EA ← (RA|0) + (RB)
if Rc = 1
    CR[CR0]LT ← 0
    CR[CR0]GT ← 0
    CR[CR0]SO ← XER[SO]
if Valid TLB entry matching EA and PID is in the TLB then
    (RT) ← Index of matching TLB Entry
    if Rc = 1
        CR[CR0]EQ ← 1
else
    (RT) Undefined
    if Rc = 1
        CR[CR0]EQ ← 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 PID. See XREF for details. The record bit (Rc) specifies whether the results of the search will affect CR[CR0] as shown above. The intention is that CR[CR0]_{EQ} can be tested after a **tlbsx.** instruction if there is a possibility that the search may fail.

Registers Altered

- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1

Invalid Instruction Forms

- None.


Programming Note

This instruction is privileged. Translation is not required to be active during the execution of this instruction.

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

tlbsync



The **tlbsync** instruction is provided in the PowerPC architecture to support synchronization of TLB operations among the processors of a multi-processor system. In the PPC405 core, this instruction performs no operation, and is provided to facilitate code portability.

Registers Altered

- None.

Invalid Instruction Forms

- None.

Programming Notes

This instruction is privileged. Translation is not required to be active during the execution of this instruction.

Since the PPC405 core does not support tightly-coupled multiprocessor systems, **tlbsync** performs no operation.


Architecture Note

This instruction is part of the IBM PowerPC Embedded Operating Environment.

tlbwe

TLB Write Entry

tlbwe RS, RA, WS



```
if WS4 = 1
    TLBLO[(RA26:31)] ← (RS)
else
    TLBHI[(RA26:31)] ← (RS)
    TID of TLB[(RA26:31)] ← (PID24:31)
```

The contents of the selected TLB entry is replaced with the contents of register RS (and possibly PID).

Bits 26:31 of the contents of RA are used as an index into the TLB. If this index specifies a TLB entry that does not exist, the results are undefined.

The WS field specifies which portion (TLBHI or TLBLO) of the entry is replaced from RS. For instructions that specify TLBHI, the TID field in the TLB entry is supplied from PID_{24:31}.

If the WS field is not 0 or 1, 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 Notes

This instruction is privileged. Translation is not required to be active during the execution of this instruction.

The effects of this update are not guaranteed to be visible to the programming model until the completion of a context synchronizing operation. For example, updating a zone selection field within the TLB while in supervisor code should be followed by an **isync** instruction (or other context synchronizing operation) to guarantee that the desired translation and protection domains are used.

tlbwe writes the TLB fields from RS and the PID as follows:

```
If WS = 0 (TLBHI):
    EPN[0:21] ← RS[0:21]
    SIZE[0:2] ← RS[22:24]
    V ← RS[25]
    E ← RS[26]
    U0 ← RS[27]
    TID[0:7] ← PID[24:31]; (note that the TID is written from the PID, not RS)
```


If WS = 1 (TLBLO):
 $RPN[0:21] \leftarrow RT[0:21]$
 $EX, WR \leftarrow RS[22:23]$
 $ZSEL[0:3] \leftarrow RS[24:27]$
 $WIMG \leftarrow RS[28:31]$

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

Table 9-33. Extended Mnemonics for tlbwe

Mnemonic	Operands	Function	Other Registers Altered
tlbwehi	RS, RA	Write TLBHI portion of the selected TLB entry from RS. Write the TID register of the selected TLB entry from the PID register. $TLBHI[(RA)] \leftarrow (RS)$ $TLB[(RA)]_{TID} \leftarrow (PID_{24:31})$ <i>Extended mnemonic for</i> tlbwe RS,RA,0	
tlbwelo	RS, RA	Write TLBLO portion of the selected TLB entry from RS. $TLBLO[(RA)] \leftarrow (RS)$ <i>Extended mnemonic for</i> tlbwe RS,RA,1	

tw

Trap Word

tw TO, RA, RB

31	TO	RA	RB	4	31
0	6	11	16	21	31

if (((RA) < (RB) \wedge TO₀ = 1) \vee
((RA) > (RB) \wedge TO₁ = 1) \vee
((RA) = (RB) \wedge TO₂ = 1) \vee
((RA) $\overset{u}{<}$ (RB) \wedge TO₃ = 1) \vee
((RA) $\overset{u}{>}$ (RB) \wedge TO₄ = 1)) then TRAP (see details below)

Register RA is compared with register RB. If any comparison condition selected by the TO field is true, a TRAP occurs. The behavior of a TRAP depends upon the debug mode of the processor, as described below:

- If TRAP is not enabled as a debug event (DBCR[TDE] = 0 or DBCR[EDM, IDM] = 0,0):

TRAP causes a program interrupt. See “Program Interrupt” on page 5-20.

(SRR0) \leftarrow address of **tw** instruction
(SRR1) \leftarrow (MSR)
(ESR[PTR]) \leftarrow 1
(MSR[WE, EE, PR, DR, IR]) \leftarrow 0
PC \leftarrow EVPR_{0:15} || 0x0700

- If TRAP is enabled as an external debug event (DBCR[TDE] = 1 and DBCR[EDM] = 1):

TRAP goes to the debug stop state, to be handled by an external debugger with hardware control.

(DBSR[TIE]) \leftarrow 1

In addition, if TRAP is also enabled as an internal debug event (DBCR[IDM] = 1) and debug exceptions are disabled (MSR[DE] = 0), then report an imprecise event:

(DBSR[IDE]) \leftarrow 1
PC \leftarrow address of **tw** instruction

- If TRAP is enabled as an internal debug event and *not* an external debug event (DBCR[TDE] = 1 and DBCR[EDM, IDM] = 0,1) and debug exceptions are enabled (MSR[DE] = 1):

TRAP causes a debug interrupt. See “Debug Interrupt” on page 5-26.

(SRR2) \leftarrow address of **tw** instruction
(SRR3) \leftarrow (MSR)
(DBSR[TIE]) \leftarrow 1
(MSR[WE, EE, PR, CE, DE, DR, IR]) \leftarrow 0
PC \leftarrow EVPR_{0:15} || 0x2000

- If TRAP is enabled as an internal debug event and *not* an external debug event (DBCR[TDE] = 1 and DBCR[EDM, IDM] = 0,1) and Debug Exceptions are disabled (MSR[DE] = 0):

TRAP reports the debug event as an *imprecise* event and causes a program interrupt. See “Program Interrupt” on page 5-20.

(SRR0) ← address of **tw** instruction
 (SRR1) ← (MSR)
 (ESR[PTR]) ← 1
 (DBSR[TIE,IDE]) ← 1,1
 (MSR[WE, EE, PR, DR, IR]) ← 0
 PC ← EVPR_{0:15} || 0x0700

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 inserted into the execution stream by a debugger to implement breakpoints, and is not typically used by application code.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-34. Extended Mnemonics for tw

Mnemonic	Operands	Function	Other Registers Altered
trap		Trap unconditionally. <i>Extended mnemonic for</i> tw 31,0,0	
tweq	RA, RB	Trap if (RA) equal to (RB). <i>Extended mnemonic for</i> tw 4,RA,RB	
twge	RA, RB	Trap if (RA) greater than or equal to (RB). <i>Extended mnemonic for</i> tw 12,RA,RB	
twgt	RA, RB	Trap if (RA) greater than (RB). <i>Extended mnemonic for</i> tw 8,RA,RB	
twle	RA, RB	Trap if (RA) less than or equal to (RB). <i>Extended mnemonic for</i> tw 20,RA,RB	
twlge	RA, RB	Trap if (RA) logically greater than or equal to (RB). <i>Extended mnemonic for</i> tw 5,RA,RB	
twlgt	RA, RB	Trap if (RA) logically greater than (RB). <i>Extended mnemonic for</i> tw 1,RA,RB	


tw

Trap Word

Table 9-34. Extended Mnemonics for tw (continued)

Mnemonic	Operands	Function	Other Registers Altered
twlle	RA, RB	Trap if (RA) logically less than or equal to (RB). <i>Extended mnemonic for</i> tw 6,RA,RB	
twllt	RA, RB	Trap if (RA) logically less than (RB). <i>Extended mnemonic for</i> tw 2,RA,RB	
twlng	RA, RB	Trap if (RA) logically not greater than (RB). <i>Extended mnemonic for</i> tw 6,RA,RB	
twlnl	RA, RB	Trap if (RA) logically not less than (RB). <i>Extended mnemonic for</i> tw 5,RA,RB	
twlt	RA, RB	Trap if (RA) less than (RB). <i>Extended mnemonic for</i> tw 16,RA,RB	
twne	RA, RB	Trap if (RA) not equal to (RB). <i>Extended mnemonic for</i> tw 24,RA,RB	
twng	RA, RB	Trap if (RA) not greater than (RB). <i>Extended mnemonic for</i> tw 20,RA,RB	
twnl	RA, RB	Trap if (RA) not less than (RB). <i>Extended mnemonic for</i> tw 12,RA,RB	

twi TO, RA, IM



if (((RA) < EXTS(IM) \wedge TO₀ = 1) \vee
 ((RA) > EXTS(IM) \wedge TO₁ = 1) \vee
 ((RA) = EXTS(IM) \wedge TO₂ = 1) \vee
 ((RA) $\overset{u}{<}$ EXTS(IM) \wedge TO₃ = 1) \vee
 ((RA) $\overset{u}{>}$ EXTS(IM) \wedge TO₄ = 1)) then TRAP (see details below)

Register RA is compared with the IM field, which has been sign-extended to 32 bits. If any comparison condition selected by the TO field is true, a TRAP occurs. The behavior of a TRAP depends upon the Debug Mode of the processor, as described below:

- If TRAP is not enabled as a debug event (DBCR[TDE] = 0 or DBCR[EDM, IDM] = 0,0):

TRAP causes a program interrupt. See “Program Interrupt” on page 5-20.

(SRR0) \leftarrow address of **twi** instruction
 (SRR1) \leftarrow (MSR)
 (ESR[PTR]) \leftarrow 1
 (MSR[WE, EE, PR, DR, IR]) \leftarrow 0
 PC \leftarrow EVPR_{0:15} || 0x0700

- If TRAP is enabled as an External debug event (DBCR[TDE] = 1 and DBCR[EDM] = 1):

TRAP goes to the Debug Stop state, to be handled by an external debugger with hardware control of the PPC405 core.

(DBSR[TIE]) \leftarrow 1
 In addition, if TRAP is also enabled as an Internal debug event (DBCR[IDM] = 1)
 and Debug Exceptions are disabled (MSR[DE] = 0), then report an imprecise event:
 (DBSR[IDE]) \leftarrow 1
 PC \leftarrow address of **twi** instruction

- If TRAP is enabled as an Internal debug event and *not* an External debug event (DBCR[TDE] = 1 and DBCR[EDM, IDM] = 0,1) and Debug Exceptions are enabled (MSR[DE] = 1):

TRAP causes a Debug interrupt. See “Debug Interrupt” on page 5-26.

(SRR2) \leftarrow address of **twi** instruction
 (SRR3) \leftarrow (MSR)
 (DBSR[TIE]) \leftarrow 1
 (MSR[WE, EE, PR, CE, DE, DR, IR]) \leftarrow 0
 PC \leftarrow EVPR_{0:15} || 0x2000

- If TRAP is enabled as an Internal debug event and *not* an External debug event (DBCR[TDE] = 1 and DBCR[EDM, IDM] = 0,1) and Debug Exceptions are disabled (MSR[DE] = 0):

TRAP will report the debug event as an *imprecise* event and will cause a Program interrupt. See “Program Interrupt” on page 5-20.

twi

Trap Word Immediate

(SRR0) ← address of **twi** instruction

(SRR1) ← (MSR)

(ESR[PTR]) ← 1

(DBSR[TIE,IDE]) ← 1,1

(MSR[WE, EE, PR, DR, IR]) ← 0

PC ← EVPR_{0:15} || 0x0700

Registers Altered

- None

Programming Note

This instruction is inserted into the execution stream by a debugger to implement breakpoints, and is not typically used by application code.

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Table 9-35. Extended Mnemonics for twi

Mnemonic	Operands	Function	Other Registers Altered
tweqi	RA, IM	Trap if (RA) equal to EXTS(IM). <i>Extended mnemonic for</i> twi 4,RA,IM	
twgei	RA, IM	Trap if (RA) greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM	
twgti	RA, IM	Trap if (RA) greater than EXTS(IM). <i>Extended mnemonic for</i> twi 8,RA,IM	
twlei	RA, IM	Trap if (RA) less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM	
twlgei	RA, IM	Trap if (RA) logically greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM	
twlgti	RA, IM	Trap if (RA) logically greater than EXTS(IM). <i>Extended mnemonic for</i> twi 1,RA,IM	
twllei	RA, IM	Trap if (RA) logically less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM	
twllti	RA, IM	Trap if (RA) logically less than EXTS(IM). <i>Extended mnemonic for</i> twi 2,RA,IM	
twlngi	RA, IM	Trap if (RA) logically not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM	


Table 9-35. Extended Mnemonics for twi (continued)

Mnemonic	Operands	Function	Other Registers Altered
twlnli	RA, IM	Trap if (RA) logically not less than EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM	
twlti	RA, IM	Trap if (RA) less than EXTS(IM). <i>Extended mnemonic for</i> twi 16,RA,IM	
twnei	RA, IM	Trap if (RA) not equal to EXTS(IM). <i>Extended mnemonic for</i> twi 24,RA,IM	
twngi	RA, IM	Trap if (RA) not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM	
twnli	RA, IM	Trap if (RA) not less than EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM	

wrtee

Write External Enable

wrtee RS



$MSR[EE] \leftarrow (RS)_{16}$

The 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 Note


Execution of this instruction is privileged.

This instruction is used to provide atomic update of MSR[EE]. Typical usage is:

```
mfmsr Rn    #save EE in Rn[16]
wrteei 0    #Turn off EE
•          #Code with EE disabled
•
•
wrtee Rn    #restore EE without affecting any MSR changes that occurred in the disabled code
```

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

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 Note

Execution of this instruction is privileged.

This instruction is used to provide an atomic update of MSR[EE]. Typical usage is:

```

mfmsr Rn    #save EE in Rn[16]
wrteei 0    #Turn off EE
•          #Code with EE disabled
•
•
wrtee Rn    #restore EE without affecting any MSR changes that occurred in the disabled code

```


Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

xor

XOR

xor RA, RS, RB Rc=0
xor. RA, RS, RB Rc=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

- CR[CR0]_{LT, GT, EQ, SO} if Rc contains 1
- RA

Architecture Note

This instruction part of the IBM PowerPC Embedded Operating Environment.

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


Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

xoris

XOR Immediate Shifted

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

Architecture Note

This instruction is part of the PowerPC User Instruction Set Architecture.

Chapter 10. Register Summary

All registers contained in the PPC405 core are architected as 32-bits. Table 10-1 and Table 10-2 define the addressing required to access the registers. The pages following these tables define the bit usage within each register.

The registers are grouped into categories, based on access mode: General Purpose Registers (GPRs), Special Purpose Registers (SPRs), Time Base Registers (TBRs), the Machine State Register (MSR), the Condition Register (CR), and, in standard products, Device Control Registers (DCRs).

10.1 Reserved Registers

Any register numbers not listed in the tables which follow are *reserved*, and should be neither read nor written. These reserved register numbers may be used for additional functions in future processors.

10.2 Reserved Fields

For all registers having 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 0 to the field. When reading from a reserved field, ignore the field.

It is good coding practice 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 General Purpose Registers

The PPC405 core provides 32 General Purpose Registers (GPRs). The contents of these registers can be loaded from memory using load instructions and stored to memory using store instructions. GPRs are also addressed by all integer instructions.

Table 10-1. PPC405 General Purpose Registers

Mnemonic	Register Name	GPR Number		Access
		Decimal	Hex	
R0–R31	General Purpose Register 0–31	0–31	0x0–0x1F	Read/Write

10.4 Machine State Register and Condition Register

Because these registers are accessed using special instructions, they do not require addressing.

10.5 Special Purpose Registers

Special Purpose Registers (SPRs), which are part of the PowerPC Embedded Environment, are accessed using the **mtspr** and **mfspr** instructions. SPRs control the use of the debug facilities, timers, interrupts, storage control attributes, and other architected processor resources.

Table 10-2 shows the mnemonics, names, and numbers of the SPRs. The columns under “SPRN” list the register numbers used as operands in assembler language coding of the **mfspr** and **mtspr** instructions. The column labeled “SPRF” lists the corresponding fields contained in the *machine code* of **mfspr** and **mtspr**. The SPRN field contains the five-bit subfields of the SPRF field, which are *reversed* in the machine code for the **mfspr** and **mtspr** instructions ($SPRN \leftarrow SPRF_{5:9} \parallel SPRF_{0:4}$) for compatibility with the POWER Architecture. Note that the assembler handles the special coding transparently.

All SPRs are privileged, except the Count Register (CTR), the Link Register (LR), SPR General Purpose Registers (SPRG4–SPRG7, read-only), User SPR General Purpose Register (USPRG0), and the Fixed-point Exception Register (XER). Note that access to the Time Base Lower (TBL) and Time Base Upper (TBU) registers, when addressed as SPRs, is write-only and privileged. However, when addressed as Time Base Registers (TBRs), read access to these registers is not privileged. See “Time Base Registers” on page 4. for more information.

Table 10-2 lists the SPRs, their mnemonics and names, their numbers (SPRN) and the corresponding SPRF numbers, and access. All SPR numbers not listed are reserved, and should be neither read nor written.

Table 10-2. Special Purpose Registers

Mnemonic	Register Name	SPRN		SPRF	Access
		Decimal	Hex		
CCR0	Core Configuration Register 0	947	0x3B3	0x27D	Read/Write
CTR	Count Register	9	0x009	0x120	Read/Write
DAC1	Data Address Compare 1	1014	0x3F6	0x2DF	Read/Write
DAC2	Data Address Compare 2	1015	0x3F7	0x2FF	Read/Write
DBCR0	Debug Control Register 0	1010	0x3F2	0x25F	Read/Write
DBCR1	Debug Control Register 1	957	0x3BD	0x3BD	Read/Write
DBSR	Debug Status Register	1008	0x3F0	0x21F	Read/Clear
DCCR	Data Cache Cachability Register	1018	0x3FA	0x35F	Read/Write
DCWR	Data Cache Write-through Register	954	0x3BA	0x35D	Read/Write
DVC1	Data Value Compare 1	950	0x3B6	0x2DD	Read/Write
DVC2	Data Value Compare 2	951	0x3B7	0x2FD	Read/Write
DEAR	Data Error Address Register	981	0x3D5	0x2BE	Read/Write
ESR	Exception Syndrome Register	980	0x3D4	0x29E	Read/Write
EVPR	Exception Vector Prefix Register	982	0x3D6	0x2DE	Read/Write
IAC1	Instruction Address Compare 1	1012	0x3F4	0x29F	Read/Write
IAC2	Instruction Address Compare 2	1013	0x3F5	0x2B5	Read/Write
IAC3	Instruction Address Compare 3	948	0x3B4	0x29D	Read/Write
IAC4	Instruction Address Compare 4	949	0x3B5	0x2BD	Read/Write

Table 10-2. Special Purpose Registers (continued)

Mnemonic	Register Name	SPRN		SPRF	Access
		Decimal	Hex		
ICCR	Instruction Cache Cachability Register	1019	0x3FB	0x37F	Read/Write
ICDBDR	Instruction Cache Debug Data Register	979	0x3D3	0x27E	Read-only
LR	Link Register	8	0x008	0x100	Read/Write
PID	Process ID	945	0x3B1	0x23D	Read/Write
PIT	Programmable Interval Timer	987	0x3DB	0x37E	Read/Write
PVR	Processor Version Register	287	0x11F	0x3E8	Read-only
SGR	Storage Guarded Register	953	0x3B9	0x33D	Read/Write
SLER	Storage Little Endian Register	955	0x3BB	0x37D	Read/Write
SPRG0	SPR General 0	272	0x110	0x208	Read/Write
SPRG1	SPR General 1	273	0x111	0x228	Read/Write
SPRG2	SPR General 2	274	0x112	0x248	Read/Write
SPRG3	SPR General 3	275	0x113	0x268	Read/Write
SPRG4	SPR General 4	260	0x104	0x088	Read-only
SPRG4	SPR General 4	276	0x114	0x288	Read/Write
SPRG5	SPR General 5	261	0x105	0x0A8	Read-only
SPRG5	SPR General 5	277	0x115	0x2A8	Read/Write
SPRG6	SPR General 6	262	0x106	0x0C8	Read-only
SPRG6	SPR General 6	278	0x116	0x2C8	Read/Write
SPRG7	SPR General 7	263	0x107	0x0E8	Read-only
SPRG7	SPR General 7	279	0x117	0x2E8	Read/Write
SRR0	Save/Restore Register 0	26	0x01A	0x340	Read/Write
SRR1	Save/Restore Register 1	27	0x01B	0x360	Read/Write
SRR2	Save/Restore Register 2	990	0x3DE	0x3DE	Read/Write
SRR3	Save/Restore Register 3	991	0x3DF	0x3FE	Read/Write
SU0R	Storage User-defined 0 Register	956	0x3BC	0x39D	Read/Write
TBL	Time Base Lower	284	0x11C	0x388	Write-only
TBU	Time Base Upper	285	0x11D	0x3A8	Write-only
TCR	Timer Control Register	986	0x3DA	0x35E	Read/Write
TSR	Timer Status Register	984	0x3D8	0x31E	Read/Clear
USPRG0	User SPR General 0	256	0x100	0x008	Read/Write
XER	Fixed Point Exception Register	1	0x001	0x020	Read/Write
ZPR	Zone Protection Register	944	0x3B0	0x21D	Privileged

10.6 Time Base Registers

The PowerPC Architecture provides a 64-bit time base. Chapter 6, “Timer Facilities,” describes the architected time base. In the PPC405 core, the time base is implemented as two 32-bit time base registers (TBRs). The low-order 32 bits of the time base are read from the TBL and the high-order 32 bits are read from the TBU.

User-mode access to the TBRs is read-only, and there is no explicitly privileged read access to the time base.

The **mftb** instruction reads from TBL and TBU. (Writing the time base is accomplished by moving the contents of a GPR to a pair of SPRs, which are also called TBL and TBU, using the **mtspr** instruction.)

Table 10-3 shows the mnemonics, names, and numbers of the TBRs. The columns under “TBRN” list the register numbers used as operands in assembler language coding of the **mftb** and **mtspr** instructions. The column labeled “TBRF” lists the corresponding fields contained in the *machine code* of **mftb** and **mtspr**. The TBRN field contains two five-bit subfields of the TBRF field; the subfields are *reversed* in the machine code for the **mftb** and **mtspr** instructions ($TBRN \leftarrow TBRF_{5:9} \parallel TBRF_{0:4}$). Note that the assembler handles the special coding transparently.

Table 10-3. Time Base Registers

Mnemonic	Register Name	TBRN		TBRF	Access
		Decimal	Hex		
TBL	Time Base Lower (Read-only)	268	0x10C	0x188	Read-only
TBU	Time Base Upper (Read-only)	269	0x10D	0x1A8	Read-only

10.7 Device Control Registers

Device Control Registers (DCRs), which are architecturally outside of the processor core, are accessed using the **mfocr** and **mtocr** 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 PPC405 core does not contain DCRs, the **mfocr** and **mtocr** 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.

10.8 Alphabetical Listing of PPC405 Registers

The following pages list the registers available in the PPC405 core. For each register, the following information is supplied:

- Register mnemonic and name
- Cross-reference to a detailed register description
- Register type (SPR or TBR; the names of CR, GPR0–31, and MSR are the same as their register types)
- Register number (address)
- 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 mnemonic, field bit location, field name, and the function associated with various field values

CCR0

Core Configuration Register 0

SPR 0x3B3

See “Core Configuration Register 0 (CCR0)” on page 4-11.




Figure 10-1. Core Configuration Register 0 (CCR0)

0:5		Reserved
6	LWL	Load Word as Line 0 The DCU performs load misses or non-cachable loads as words, halfwords, or bytes, as requested 1 For load misses or non-cachable loads, the DCU moves eight words (including the target word) into the line fill buffer
7	LWOA	Load Without Allocate 0 Load misses result in line fills 1 Load misses do not result in a line fill, but in non-cachable loads
8	SWOA	Store Without Allocate 0 Store misses result in line fills 1 Store misses do not result in line fills, but in non-cachable stores
9	DPP1	DCU PLB Priority Bit 1 0 DCU PLB priority 0 on bit 1 1 DCU PLB priority 1 on bit 1 Note: DCU logic dynamically controls DCU priority bit 0.
10:11	IPP	ICU PLB Priority Bits 0:1 00 Lowest ICU PLB priority 01 Next to lowest ICU PLB priority 10 Next to highest ICU PLB priority 11 Highest ICU PLB priority
12:13		Reserved
14	U0XE	Enable U0 Exception 0 Disables the U0 exception 1 Enables the U0 exception
15	LBDE	Load Debug Enable 0 Load data is invisible on data-side (on-chip memory (OCM)) 1 Load data is visible on data-side OCM
16:19		Reserved
20	PFC	ICU Prefetching for Cachable Regions 0 Disables prefetching for cachable regions 1 Enables prefetching for cachable regions

CCR0 (cont.)

Core Configuration Register 0

21	PFNC	ICU Prefetching for Non-Cachable Regions 0 Disables prefetching for non-cachable regions 1 Enables prefetching for non-cachable regions
22	NCRS	Non-cachable ICU request size 0 Requests are for four-word lines 1 Requests are for eight-word lines
23	FWOA	Fetch Without Allocate 0 An ICU miss results in a line fill. 1 An ICU miss does not cause a line fill, but results in a non-cachable fetch.
24:26		Reserved
27	CIS	Cache Information Select 0 Information is cache data. 1 Information is cache tag.
28:30		Reserved
31	CWS	Cache Way Select 0 Cache way is A. 1 Cache way is B.

CR

Condition Register

See “Condition Register (CR)” on page 2-10.




Figure 10-2. 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

SPR 0x009

See “Count Register (CTR)” on page 2-6.

**Figure 10-3. Count Register (CTR)**

0:31	Count	Used as count for branch conditional with decrement instructions, or as address for branch-to-counter instructions.
------	-------	---

DAC1–DAC2

Data Address Compare Registers

SPR 0x3F6–0x3F7

See “Data Address Compare Registers (DAC1–DAC2)” on page 8-9.



Figure 10-4. Data Address Compare Registers (DAC1–DAC2)

0:31		Data Address Compare (DAC) byte address	DBCRO[D1S] determines which address bits are examined.
------	--	---	--

SPR 0x3F2

See “Debug Control Registers” on page 8-4.




Figure 10-5. Debug Control Register 0 (DBCRO)

0	EDM	External Debug Mode 0 Disabled 1 Enabled	
1	IDM	Internal Debug Mode 0 Disabled 1 Enabled	
2:3	RST	Reset 00 No action 01 Core reset 10 Chip reset 11 System reset	Causes a processor reset request when set by software.
Attention: Writing 01, 10, or 11 to this field causes a processor reset request.			
4	IC	Instruction Completion Debug Event 0 Disabled 1 Enabled	
5	BT	Branch Taken Debug Event 0 Disabled 1 Enabled	
6	EDE	Exception Debug Event 0 Disabled 1 Enabled	
7	TDE	Trap Debug Event 0 Disabled 1 Enabled	
8	IA1	IAC 1 Debug Event 0 Disabled 1 Enabled	
9	IA2	IAC 2 Debug Event 0 Disabled 1 Enabled	
10	IA12	Instruction Address Range Compare 1-2 0 Disabled 1 Enabled	Registers IAC1 and IAC2 define an address range used for IAC address comparisons.

DBCR0 (cont.)

Debug Control Register 0

11	IA12X	Enable Instruction Address Exclusive Range Compare 1–2 0 Inclusive 1 Exclusive	Selects the range defined by IAC1 and IAC2 to be inclusive or exclusive.
12	IA3	IAC 3 Debug Event 0 Disabled 1 Enabled	
13	IA4	IAC 4 Debug Event 0 Disabled 1 Enabled	
14	IA34	Instruction Address Range Compare 3–4 0 Disabled 1 Enabled	Registers IAC3 and IAC4 define an address range used for IAC address comparisons.
15	IA34X	Instruction Address Exclusive Range Compare 3–4 0 Inclusive 1 Exclusive	Selects range defined by IAC3 and IAC4 to be inclusive or exclusive.
16	IA12T	Instruction Address Range Compare 1-2 Toggle 0 Disabled 1 Enable	Toggles range 12 inclusive, exclusive DBCR[IA12X] on debug event.
17	IA34T	Instruction Address Range Compare 3–4 Toggle 0 Disabled 1 Enable	Toggles range 34 inclusive, exclusive DBCR[IA34X] on debug event.
18:30		Reserved	
31	FT	Freeze timers on debug event 0 Timers not frozen 1 Timers frozen	

SPR 0x3BD

See “Debug Control Registers” on page 8-4.




Figure 10-6. Debug Control Register 1 (DBCR1)

0	D1R	DAC1 Read Debug Event 0 Disabled 1 Enabled	
1	D2R	DAC 2 Read Debug Event 0 Disabled 1 Enabled	
2	D1W	DAC 1 Write Debug Event 0 Disabled 1 Enabled	
3	D2W	DAC 2 Write Debug Event 0 Disabled 1 Enabled	
4:5	D1S	DAC 1 Size 00 Compare all bits 01 Ignore lsb (least significant bit) 10 Ignore two lsbs 11 Ignore five lsbs	Address bits used in the compare: Byte address Halfword address Word address Cache line (8-word) address
6:7	D2S	DAC 2 Size 00 Compare all bits 01 Ignore lsb (least significant bit) 10 Ignore two lsbs 11 Ignore five lsbs	Address bits used in the compare: Byte address Halfword address Word address Cache line (8-word) address
8	DA12	Enable Data Address Range Compare 1:2 0 Disabled 1 Enabled	Registers DAC1 and DAC2 define an address range used for DAC address comparisons
9	DA12X	Data Address Exclusive Range Compare 1:2 0 Inclusive 1 Exclusive	Selects range defined by DAC1 and DAC2 to be inclusive or exclusive
10:11		Reserved	

DBCR1 (cont.)

Debug Control Register 1

12:13	DV1M	Data Value Compare 1 Mode 00 Undefined 01 AND 10 OR 11 AND-OR	Type of data comparison used: All bytes selected by DBCR1[DV1BE] must compare to the appropriate bytes of DVC1. One of the bytes selected by DBCR1[DV1BE] must compare to the appropriate bytes of DVC1. The upper halfword or lower halfword must compare to the appropriate halfword in DVC1. When performing halfword compares set DBCR1[DV1BE] = 0011, 1100, or 1111.
14:15	DV2M	Data Value Compare 2 Mode 00 Undefined 01 AND 10 OR 11 AND-OR	Type of data comparison used All bytes selected by DBCR1[DV2BE] must compare to the appropriate bytes of DVC2. One of the bytes selected by DBCR1[DV2BE] must compare to the appropriate bytes of DVC2. The upper halfword or lower halfword must compare to the appropriate halfword in DVC2. When performing halfword compares set DBCR1[DV2BE] = 0011, 1100, or 1111.
16:19	DV1BE	Data Value Compare 1 Byte 0 Disabled 1 Enabled	Selects which data bytes to use in data value comparison
20:23	DV2BE	Data Value Compare 2 Byte 0 Disabled 1 Enabled	Selects which data bytes to use in data value comparison
24:31		Reserved	

SPR 0x3F0 Read/Clear

See “Debug Status Register (DBSR)” on page 8-7.




Figure 10-7. Debug Status Register (DBSR)

0	IC	Instruction Completion Debug Event 0 Event did not occur 1 Event occurred
1	BT	Branch Taken Debug Event 0 Event did not occur 1 Event occurred
2	EDE	Exception Debug Event 0 Event did not occur 1 Event occurred
3	TIE	Trap Instruction Debug Event 0 Event did not occur 1 Event occurred
4	UDE	Unconditional Debug Event 0 Event did not occur 1 Event occurred
5	IA1	IAC1 Debug Event 0 Event did not occur 1 Event occurred
6	IA2	IAC2 Debug Event 0 Event did not occur 1 Event occurred
7	DR1	DAC1 Read Debug Event 0 Event did not occur 1 Event occurred
8	DW1	DAC1 Write Debug Event 0 Event did not occur 1 Event occurred
9	DR2	DAC2 Read Debug Event 0 Event did not occur 1 Event occurred
10	DW2	DAC2 Write Debug Event 0 Event did not occur 1 Event occurred

DBSR (cont.)

Debug Status Register

11	IDE	Imprecise Debug Event 0 No circumstance that would cause a debug event (if MSR[DE] = 1) occurred 1 A debug event would have occurred, but debug exceptions were disabled (MSR[DE] = 0)	
12	IA3	IAC3 Debug Event 0 Event did not occur 1 Event occurred	
13	IA4	IAC4 Debug Event 0 Event did not occur 1 Event occurred	
14:21		Reserved	
22:23	MRR	Most Recent Reset 00 No reset has occurred since last cleared by software. 01 Core reset 10 Chip reset 11 System reset	This field is set to a value, indicating the type of reset, when a reset occurs.
24:31		Reserved	

SPR 0x3FA

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-8. Data Cache Cachability Register (DCCR)

0	S0	0 Noncachable 1 Cachable	0x0000 0000–0x07FF FFFF
1	S1	0 Noncachable 1 Cachable	0x0800 0000–0x0FFF FFFF
2	S2	0 Noncachable 1 Cachable	0x1000 0000–0x17FF FFFF
3	S3	0 Noncachable 1 Cachable	0x1800 0000–0x1FFF FFFF
4	S4	0 Noncachable 1 Cachable	0x2000 0000–0x27FF FFFF
5	S5	0 Noncachable 1 Cachable	0x2800 0000–0x2FFF FFFF
6	S6	0 Noncachable 1 Cachable	0x3000 0000–0x37FF FFFF
7	S7	0 Noncachable 1 Cachable	0x3800 0000–0x3FFF FFFF
8	S8	0 Noncachable 1 Cachable	0x4000 0000–0x47FF FFFF
9	S9	0 Noncachable 1 Cachable	0x4800 0000–0x4FFF FFFF
10	S10	0 Noncachable 1 Cachable	0x5000 0000–0x57FF FFFF
11	S11	0 Noncachable 1 Cachable	0x5800 0000–0x5FFF FFFF
12	S12	0 Noncachable 1 Cachable	0x6000 0000–0x67FF FFFF
13	S13	0 Noncachable 1 Cachable	0x6800 0000–0x6FFF FFFF
14	S14	0 Noncachable 1 Cachable	0x7000 0000–0x77FF FFFF
15	S15	0 Noncachable 1 Cachable	0x7800 0000–0x7FFF FFFF

DCCR (cont.)

Data Cache Cacheability Register

16	S16	0 Noncachable 1 Cachable	0x8000 0000–0x87FF FFFF
17	S17	0 Noncachable 1 Cachable	0x8800 0000–0x8FFF FFFF
18	S18	0 Noncachable 1 Cachable	0x9000 0000–0x97FF FFFF
19	S19	0 Noncachable 1 Cachable	0x9800 0000–0x9FFF FFFF
20	S20	0 Noncachable 1 Cachable	0xA000 0000–0xA7FF FFFF
21	S21	0 Noncachable 1 Cachable	0xA800 0000–0xAFFF FFFF
22	S22	0 Noncachable 1 Cachable	0xB000 0000–0xB7FF FFFF
23	S23	0 Noncachable 1 Cachable	0xB800 0000–0xBFFF FFFF
24	S24	0 Noncachable 1 Cachable	0xC000 0000–0xC7FF FFFF
25	S25	0 Noncachable 1 Cachable	0xC800 0000–0xCFFF FFFF
26	S26	0 Noncachable 1 Cachable	0xD000 0000–0xD7FF FFFF
27	S27	0 Noncachable 1 Cachable	0xD800 0000–0xDFFF FFFF
28	S28	0 Noncachable 1 Cachable	0xE000 0000–0xE7FF FFFF
29	S29	0 Noncachable 1 Cachable	0xE800 0000–0xEFFF FFFF
30	S30	0 Noncachable 1 Cachable	0xF000 0000–0xF7FF FFFF
31	S31	0 Noncachable 1 Cachable	0xF800 0000–0xFFFF FFFF

SPR 0x3BA

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-9. Data Cache Write-through Register (DCWR)

0	W0	0 Write-back 1 Write-through	0x0000 0000–0x07FF FFFF
1	W1	0 Write-back 1 Write-through	0x0800 0000–0x0FFF FFFF
2	W2	0 Write-back 1 Write-through	0x1000 0000–0x17FF FFFF
3	W3	0 Write-back 1 Write-through	0x1800 0000–0x1FFF FFFF
4	W4	0 Write-back 1 Write-through	0x2000 0000–0x27FF FFFF
5	W5	0 Write-back 1 Write-through	0x2800 0000–0x2FFF FFFF
6	W6	0 Write-back 1 Write-through	0x3000 0000–0x37FF FFFF
7	W7	0 Write-back 1 Write-through	0x3800 0000–0x3FFF FFFF
8	W8	0 Write-back 1 Write-through	0x4000 0000–0x47FF FFFF
9	W9	0 Write-back 1 Write-through	0x4800 0000–0x4FFF FFFF
10	W10	0 Write-back 1 Write-through	0x5000 0000–0x57FF FFFF
11	W11	0 Write-back 1 Write-through	0x5800 0000–0x5FFF FFFF
12	W12	0 Write-back 1 Write-through	0x6000 0000–0x67FF FFFF
13	W13	0 Write-back 1 Write-through	0x6800 0000–0x6FFF FFFF
14	W14	0 Write-back 1 Write-through	0x7000 0000–0x77FF FFFF
15	W15	0 Write-back 1 Write-through	0x7800 0000–0x7FFF FFFF

DCWR (cont.)

Data Cache Write-through Register

16	W16	0 Write-back 1 Write-through	0x8000 0000–0x87FF FFFF
17	W17	0 Write-back 1 Write-through	0x8800 0000–0x8FFF FFFF
18	W18	0 Write-back 1 Write-through	0x9000 0000–0x97FF FFFF
19	W19	0 Write-back 1 Write-through	0x9800 0000–0x9FFF FFFF
20	W20	0 Write-back 1 Write-through	0xA000 0000–0xA7FF FFFF
21	W21	0 Write-back 1 Write-through	0xA800 0000–0xAFFF FFFF
22	W22	0 Write-back 1 Write-through	0xB000 0000–0xB7FF FFFF
23	W23	0 Write-back 1 Write-through	0xB800 0000–0xBFFF FFFF
24	W24	0 Write-back 1 Write-through	0xC000 0000–0xC7FF FFFF
25	W25	0 Write-back 1 Write-through	0xC800 0000–0xCFFF FFFF
26	W26	0 Write-back 1 Write-through	0xD000 0000–0xD7FF FFFF
27	W27	0 Write-back 1 Write-through	0xD800 0000–0xDFFF FFFF
28	W28	0 Write-back 1 Write-through	0xE000 0000–0xE7FF FFFF
29	W29	0 Write-back 1 Write-through	0xE800 0000–0xEFFF FFFF
30	W30	0 Write-back 1 Write-through	0xF000 0000–0xF7FF FFFF
31	W31	0 Write-back 1 Write-through	0xF800 0000–0xFFFF FFFF

DEAR


Data Exception Address Register

SPR 0x3D5

See “Data Exception Address Register (DEAR)” on page 5-13.



Figure 10-10. Data Exception Address Register (DEAR)



DVC1–DVC2

Data Value Compare Registers

SPR 0x3B6–0x3B7

See “Data Value Compare Registers (DVC1–DVC2)” on page 8-10.




Figure 10-11. Data Value Compare Registers (DVC1–DVC2)



SPR 0x3D4

See “Exception Syndrome Register (ESR)” on page 5-11.




Figure 10-12. Exception Syndrome Register (ESR)

0	MCI	Machine check—instruction 0 Instruction machine check did not occur. 1 Instruction machine check occurred.
1:3		Reserved
4	PIL	Program interrupt—illegal 0 Illegal Instruction error did not occur. 1 Illegal Instruction error occurred.
5	PPR	Program interrupt—privileged 0 Privileged instruction error did not occur. 1 Privileged instruction error occurred.
6	PTR	Program interrupt—trap 0 Trap with successful compare did not occur. 1 Trap with successful compare occurred.
7	PEU	Program interrupt—Unimplemented 0 APU/FPU unimplemented exception did not occur. 1 APU/FPU unimplemented exception occurred.
8	DST	Data storage interrupt—store fault 0 Excepting instruction was not a store. 1 Excepting instruction was a store (includes dcbi , dcbz , and dccci).
9	DIZ	Data/instruction storage interrupt—zone fault 0 Excepting condition was not a zone fault. 1 Excepting condition was a zone fault.
10:11		Reserved
12	PFP	Program interrupt—FPU 0 FPU interrupt did not occur. 1 FPU interrupt occurred.
13	PAP	Program interrupt—APU 0 APU interrupt did not occur. 1 APU interrupt occurred.
14:15		Reserved

ESR (cont.)

Exception Syndrome Register

16	U0F	Data storage interrupt—U0 fault 0 Excepting instruction did not cause a U0 fault. 1 Excepting instruction did cause a U0 fault.
17:31		Reserved

EVPR

Exception Vector Prefix Register

SPR 0x3D6

See “Exception Vector Prefix Register (EVPR)” on page 5-10.




Figure 10-13. Exception Vector Prefix Register (EVPR)

0:15	EVP	Exception Vector Prefix
16:31		Reserved

GPR0–GPR31

General Purpose Registers

See “General Purpose Registers (R0-R31)” on page 2-5.



Figure 10-14. General Purpose Registers (R0-R31)



IAC1–IAC4

Instruction Address Compare Registers

SPR 0x3F4–0x3F5

See “Instruction Address Compare Registers (IAC1–IAC4)” on page 8-9.



Figure 10-15. Instruction Address Compare Registers (IAC1–IAC4)

0:29		Instruction Address Compare word address	Omit two low-order bits of complete address.
30:31		Reserved	

ICCR

Instruction Cache Cacheability Register

SPR 0x3FB

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-16. Instruction Cache Cachability Register (ICCR)

0	S0	0 Noncachable 1 Cachable	0x0000 0000–0x07FF FFFF
1	S1	0 Noncachable 1 Cachable	0x0800 0000–0x0FFF FFFF
2	S2	0 Noncachable 1 Cachable	0x1000 0000–0x17FF FFFF
3	S3	0 Noncachable 1 Cachable	0x1800 0000–0x1FFF FFFF
4	S4	0 Noncachable 1 Cachable	0x2000 0000–0x27FF FFFF
5	S5	0 Noncachable 1 Cachable	0x2800 0000–0x2FFF FFFF
6	S6	0 Noncachable 1 Cachable	0x3000 0000–0x37FF FFFF
7	S7	0 Noncachable 1 Cachable	0x3800 0000–0x3FFF FFFF
8	S8	0 Noncachable 1 Cachable	0x4000 0000–0x47FF FFFF
9	S9	0 Noncachable 1 Cachable	0x4800 0000–0x4FFF FFFF
10	S10	0 Noncachable 1 Cachable	0x5000 0000–0x57FF FFFF
11	S11	0 Noncachable 1 Cachable	0x5800 0000–0x5FFF FFFF
12	S12	0 Noncachable 1 Cachable	0x6000 0000–0x67FF FFFF
13	S13	0 Noncachable 1 Cachable	0x6800 0000–0x6FFF FFFF
14	S14	0 Noncachable 1 Cachable	0x7000 0000–0x77FF FFFF
15	S15	0 Noncachable 1 Cachable	0x7800 0000–0x7FFF FFFF

ICCR (cont.)

Instruction Cache Cacheability Register

16	S16	0 Noncachable 1 Cachable	0x8000 0000–0x87FF FFFF
17	S17	0 Noncachable 1 Cachable	0x8800 0000–0x8FFF FFFF
18	S18	0 Noncachable 1 Cachable	0x9000 0000–0x97FF FFFF
19	S19	0 Noncachable 1 Cachable	0x9800 0000–0x9FFF FFFF
20	S20	0 Noncachable 1 Cachable	0xA000 0000–0xA7FF FFFF
21	S21	0 Noncachable 1 Cachable	0xA800 0000–0xAFFF FFFF
22	S22	0 Noncachable 1 Cachable	0xB000 0000–0xB7FF FFFF
23	S23	0 Noncachable 1 Cachable	0xB800 0000–0xBFFF FFFF
24	S24	0 Noncachable 1 Cachable	0xC000 0000–0xC7FF FFFF
25	S25	0 Noncachable 1 Cachable	0xC800 0000–0xCFFF FFFF
26	S26	0 Noncachable 1 Cachable	0xD000 0000–0xD7FF FFFF
27	S27	0 Noncachable 1 Cachable	0xD800 0000–0xDFFF FFFF
28	S28	0 Noncachable 1 Cachable	0xE000 0000–0xE7FF FFFF
29	S29	0 Noncachable 1 Cachable	0xE800 0000–0xEFFF FFFF
30	S30	0 Noncachable 1 Cachable	0xF000 0000–0xF7FF FFFF
31	S31	0 Noncachable 1 Cachable	0xF800 0000–0xFFFF FFFF

ICDBDR

Instruction Cache Debug Data Register

SPR 0x3D3 Read-Only

See “ICU Debugging” on page 4-14.

0	31
---	----

Figure 10-17. Instruction Cache Debug Data Register (ICDBDR)

0:31	Instruction cache information	See icread , page -68.
------	-------------------------------	-------------------------------

ICU tag information is placed into the ICDBDR as shown:

0:21	TAG	Cache Tag
22:26		Reserved
27	V	Cache Line Valid 0 Not valid 1 Valid
28:30		Reserved
31	LRU	Least Recently Used (LRU) 0 A-way LRU 1 B-way LRU

SPR 0x008

See “Link Register (LR)” on page 2-7.

**Figure 10-18. Link Register (LR)**

0:31	Link Register contents	If (LR) represents an instruction address, LR _{30:31} should be 0.
------	------------------------	--

MSR

Machine State Register

See “Machine State Register (MSR)” on page 5-7.




Figure 10-19. Machine State Register (MSR)

0:5		Reserved	
6	AP	Auxiliary Processor Available 0 APU not available. 1 APU available.	
7:11		Reserved	
12	APE	APU Exception Enable 0 APU exception disabled. 1 APU exception enabled.	
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 interrupts are disabled. 1 Critical interrupts are enabled.	Controls the critical interrupt input and watchdog timer first time-out interrupts.
15		Reserved	
16	EE	External Interrupt Enable 0 Asynchronous interrupts are disabled. 1 Asynchronous interrupts are enabled.	Controls the non-critical external interrupt input, PIT, and FIT interrupts.
17	PR	Problem State 0 Supervisor state (all instructions allowed). 1 Problem state (some instructions not allowed).	
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.	

MSR (cont.)

Machine State Register

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 Debug wait mode is disabled. 1 Debug wait mode is enabled.
22	DE	Debug Interrupts 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	IR	Instruction Relocate 0 Instruction address translation is disabled. 1 Instruction address translation is enabled.
27	DR	Data Relocate 0 Data address translation is disabled. 1 Data address translation is enabled.
28:31		Reserved

PID

Process ID

SPR 0x3B1

See "Address Translation" on page 7-1.



Figure 10-20. Process ID (PID)

0:23		Reserved
24:31		Process ID

SPR 0x3DB

See “Programmable Interval Timer (PIT)” on page 6-4.

0	31
---	----

Figure 10-21. Programmable Interval Timer (PIT)

0:31		Programmed interval remaining	Number of clocks remaining until the PIT event
------	--	-------------------------------	--

PVR

Processor Version Register

SPR 0x11F Read-Only

See “Processor Version Register (PVR)” on page 2-10.




Figure 10-22. Processor Version Register (PVR)

0:11	OWN	Owner Identifier	Identifies the owner of a core
12:15	PCF	Processor Core Family	Identifies the processor core family.
16:21	CAS	Cache Array Sizes	Identifies the cache array sizes.
22:25	PCL	Processor Core Version	Identifies the core version for a specific combination of PVR[PCF] and PVR[CAS]
26:31	AID	ASIC Identifier	Assigned sequentially; identifies an ASIC function, version, and technology

SPR 0x3B9

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-23. Storage Guarded Register (SGR)

0	G0	0 Normal 1 Guarded	0x0000 0000–0x07FF FFFF
1	G1	0 Normal 1 Guarded	0x0800 0000–0x0FFF FFFF
2	G2	0 Normal 1 Guarded	0x1000 0000–0x17FF FFFF
3	G3	0 Normal 1 Guarded	0x1800 0000–0x1FFF FFFF
4	G4	0 Normal 1 Guarded	0x2000 0000–0x27FF FFFF
5	G5	0 Normal 1 Guarded	0x2800 0000–0x2FFF FFFF
6	G6	0 Normal 1 Guarded	0x3000 0000–0x37FF FFFF
7	G7	0 Normal 1 Guarded	0x3800 0000–0x3FFF FFFF
8	G8	0 Normal 1 Guarded	0x4000 0000–0x47FF FFFF
9	G9	0 Normal 1 Guarded	0x4800 0000–0x4FFF FFFF
10	G10	0 Normal 1 Guarded	0x5000 0000–0x57FF FFFF
11	G11	0 Normal 1 Guarded	0x5800 0000–0x5FFF FFFF
12	G12	0 Normal 1 Guarded	0x6000 0000–0x67FF FFFF
13	G13	0 Normal 1 Guarded	0x6800 0000–0x6FFF FFFF
14	G14	0 Normal 1 Guarded	0x7000 0000–0x77FF FFFF
15	G15	0 Normal 1 Guarded	0x7800 0000–0x7FFF FFFF

SGR (cont.)

Storage Guarded Register

16	G16	0 Normal 1 Guarded	0x8000 0000–0x87FF FFFF
17	G17	0 Normal 1 Guarded	0x8800 0000–0x8FFF FFFF
18	G18	0 Normal 1 Guarded	0x9000 0000–0x97FF FFFF
19	G19	0 Normal 1 Guarded	0x9800 0000–0x9FFF FFFF
20	G20	0 Normal 1 Guarded	0xA000 0000–0xA7FF FFFF
21	G21	0 Normal 1 Guarded	0xA800 0000–0xAFFF FFFF
22	G22	0 Normal 1 Guarded	0xB000 0000–0xB7FF FFFF
23	G23	0 Normal 1 Guarded	0xB800 0000–0xBFFF FFFF
24	G24	0 Normal 1 Guarded	0xC000 0000–0xC7FF FFFF
25	G25	0 Normal 1 Guarded	0xC800 0000–0xCFFF FFFF
26	G26	0 Normal 1 Guarded	0xD000 0000–0xD7FF FFFF
27	G27	0 Normal 1 Guarded	0xD800 0000–0xDFFF FFFF
28	G28	0 Normal 1 Guarded	0xE000 0000–0xE7FF FFFF
29	G29	0 Normal 1 Guarded	0xE800 0000–0xEFFF FFFF
30	G30	0 Normal 1 Guarded	0xF000 0000–0xF7FF FFFF
31	G31	0 Normal 1 Guarded	0xF800 0000–0xFFFF FFFF

SPR 0x3BB

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-24. Storage Little-Endian Register (SLER)

0	S0	0 Big endian 1 Little endian	0x0000 0000–0x07FF FFFF
1	S1	0 Big endian 1 Little endian	0x0800 0000–0x0FFF FFFF
2	S2	0 Big endian 1 Little endian	0x1000 0000–0x17FF FFFF
3	S3	0 Big endian 1 Little endian	0x1800 0000–0x1FFF FFFF
4	S4	0 Big endian 1 Little endian	0x2000 0000–0x27FF FFFF
5	S5	0 Big endian 1 Little endian	0x2800 0000–0x2FFF FFFF
6	S6	0 Big endian 1 Little endian	0x3000 0000–0x37FF FFFF
7	S7	0 Big endian 1 Little endian	0x3800 0000–0x3FFF FFFF
8	S8	0 Big endian 1 Little endian	0x4000 0000–0x47FF FFFF
9	S9	0 Big endian 1 Little endian	0x4800 0000–0x4FFF FFFF
10	S10	0 Big endian 1 Little endian	0x5000 0000–0x57FF FFFF
11	S11	0 Big endian 1 Little endian	0x5800 0000–0x5FFF FFFF
12	S12	0 Big endian 1 Little endian	0x6000 0000–0x67FF FFFF
13	S13	0 Big endian 1 Little endian	0x6800 0000–0x6FFF FFFF
14	S14	0 Big endian 1 Little endian	0x7000 0000–0x77FF FFFF
15	S15	0 Big endian 1 Little endian	0x7800 0000–0x7FFF FFFF

SLER (cont.)

Storage Little-Endian Register

16	S16	0 Big endian 1 Little endian	0x8000 0000–0x87FF FFFF
17	S17	0 Big endian 1 Little endian	0x8800 0000–0x8FFF FFFF
18	S18	0 Big endian 1 Little endian	0x9000 0000–0x97FF FFFF
19	S19	0 Big endian 1 Little endian	0x9800 0000–0x9FFF FFFF
20	S20	0 Big endian 1 Little endian	0xA000 0000–0xA7FF FFFF
21	S21	0 Big endian 1 Little endian	0xA800 0000–0xAFFF FFFF
22	S22	0 Big endian 1 Little endian	0xB000 0000–0xB7FF FFFF
23	S23	0 Big endian 1 Little endian	0xB800 0000–0xBFFF FFFF
24	S24	0 Big endian 1 Little endian	0xC000 0000–0xC7FF FFFF
25	S25	0 Big endian 1 Little endian	0xC800 0000–0xCFFF FFFF
26	S26	0 Big endian 1 Little endian	0xD000 0000–0xD7FF FFFF
27	S27	0 Big endian 1 Little endian	0xD800 0000–0xDFFF FFFF
28	S28	0 Big endian 1 Little endian	0xE000 0000–0xE7FF FFFF
29	S29	0 Big endian 1 Little endian	0xE800 0000–0xEFFF FFFF
30	S30	0 Big endian 1 Little endian	0xF000 0000–0xF7FF FFFF
31	S31	0 Big endian 1 Little endian	0xF800 0000–0xFFFF FFFF

SPRG0–SPRG7

Special Purpose Registers General

SPR 0x104–0x107 (User Read-only); 0x110–0x117 (Privileged Read/Write)

See “Special Purpose Register General (SPRG0–SPRG7)” on page 2-9.

0	31
---	----

Figure 10-25. Special Purpose Registers General (SPRG0–SPRG7)

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

SRR0

Save/Restore Register 0

SPR 0x01A

See “Save/Restore Registers 0 and 1 (SRR0–SRR1)” on page 5-9.



Figure 10-26. Save/Restore Register 0 (SRR0)

0:29		SRR0 receives an instruction address when a non-critical interrupt is taken; the Program Counter is restored from SRR0 when rfi executes.
30:31		Reserved

SPR 0x01B

See “Save/Restore Registers 0 and 1 (SRR0–SRR1)” on page 5-9.




Figure 10-27. Save/Restore Register 1 (SRR1)

0:31	SRR1 receives a copy of the MSR when an interrupt is taken; the MSR is restored from SRR1 when <code>rfi</code> executes.
------	---

SRR2

Save/Restore Register 2

SPR 0x3DE

See “Save/Restore Registers 2 and 3 (SRR2–SRR3)” on page 5-9.



Figure 10-28. Save/Restore Register 2 (SRR2)

0:29		SRR2 receives an instruction address when a critical interrupt is taken; the Program Counter is restored from SRR2 when rftci executes.
30:31		Reserved

SRR3

Save/Restore Register 3

SPR 0x3DF

See “Save/Restore Registers 2 and 3 (SRR2–SRR3)” on page 5-9.




Figure 10-29. Save/Restore Register 3 (SRR3)

0:31	SRR3 receives a copy of the MSR when a critical interrupt is taken; the MSR is restored from SRR3 when rfci executes.
------	--

SU0R

Storage User-Defined 0 Register

SPR 0x3BC

See “Real-Mode Storage Attribute Control” on page 7-17.




Figure 10-30. Storage User-defined 0 Register (SU0R)

0	UD0	0 Storage compression is off 1 Storage compression is on	0x0000 0000–0x07FF FFFF
1	UD1	0 Storage compression is off 1 Storage compression is on	0x0800 0000–0x0FFF FFFF
2	UD2	0 Storage compression is off 1 Storage compression is on	0x1000 0000–0x17FF FFFF
3	UD3	0 Storage compression is off 1 Storage compression is on	0x1800 0000–0x1FFF FFFF
4	UD4	0 Storage compression is off 1 Storage compression is on	0x2000 0000–0x27FF FFFF
5	UD5	0 Storage compression is off 1 Storage compression is on	0x2800 0000–0x2FFF FFFF
6	UD6	0 Storage compression is off 1 Storage compression is on	0x3000 0000–0x37FF FFFF
7	UD7	0 Storage compression is off 1 Storage compression is on	0x3800 0000–0x3FFF FFFF
8	UD8	0 Storage compression is off 1 Storage compression is on	0x4000 0000–0x47FF FFFF
9	UD9	0 Storage compression is off 1 Storage compression is on	0x4800 0000–0x4FFF FFFF
10	UD10	0 Storage compression is off 1 Storage compression is on	0x5000 0000–0x57FF FFFF
11	UD11	0 Storage compression is off 1 Storage compression is on	0x5800 0000–0x5FFF FFFF
12	UD12	0 Storage compression is off 1 Storage compression is on	0x6000 0000–0x67FF FFFF
13	UD13	0 Storage compression is off 1 Storage compression is on	0x6800 0000–0x6FFF FFFF
14	UD14	0 Storage compression is off 1 Storage compression is on	0x7000 0000–0x77FF FFFF
15	UD15	0 Storage compression is off 1 Storage compression is on	0x7800 0000–0x7FFF FFFF

SU0R (cont.)

Storage User-Defined 0 Register

16	UD16	0 Storage compression is off 1 Storage compression is on	0x8000 0000–0x87FF FFFF
17	UD17	0 Storage compression is off 1 Storage compression is on	0x8800 0000–0x8FFF FFFF
18	UD18	0 Storage compression is off 1 Storage compression is on	0x9000 0000–0x97FF FFFF
19	UD19	0 Storage compression is off 1 Storage compression is on	0x9800 0000–0x9FFF FFFF
20	UD20	0 Storage compression is off 1 Storage compression is on	0xA000 0000–0xA7FF FFFF
21	UD21	0 Storage compression is off 1 Storage compression is on	0xA800 0000–0xAFFF FFFF
22	UD22	0 Storage compression is off 1 Storage compression is on	0xB000 0000–0xB7FF FFFF
23	UD23	0 Storage compression is off 1 Storage compression is on	0xB800 0000–0xBFFF FFFF
24	UD24	0 Storage compression is off 1 Storage compression is on	0xC000 0000–0xC7FF FFFF
25	UD25	0 Storage compression is off 1 Storage compression is on	0xC800 0000–0xCFFF FFFF
26	UD26	0 Storage compression is off 1 Storage compression is on	0xD000 0000–0xD7FF FFFF
27	UD27	0 Storage compression is off 1 Storage compression is on	0xD800 0000–0xDFFF FFFF
28	UD28	0 Storage compression is off 1 Storage compression is on	0xE000 0000–0xE7FF FFFF
29	UD29	0 Storage compression is off 1 Storage compression is on	0xE800 0000–0xEFFF FFFF
30	UD30	0 Storage compression is off 1 Storage compression is on	0xF000 0000–0xF7FF FFFF
31	UD31	0 Storage compression is off 1 Storage compression is on	0xF800 0000–0xFFFF FFFF

TBL

Time Base Lower

TBR 0x10C (Read-only); SPR 0x11C (Privileged write-only)

See “Time Base” on page 6-1.



Figure 10-31. Time Base Lower (TBL)

0:31		Time Base Lower	Current count; low-order 32 bits of time base.
------	--	-----------------	--

TBR 0x10D (Read-only); SPR 0x11D (Privileged write-only)

See “Time Base” on page 6-1.



Figure 10-32. Time Base Upper (TBU)

0:31	Time Base Upper	Current count, high-order 32 bits of time base.
------	-----------------	---

TCR

Timer Control Register

SPR 0x3DA

See “Timer Control Register (TCR)” on page 6-9.




Figure 10-33. Timer Control Register (TCR)

0:1	WP	Watchdog Period 00 2^{17} clocks 01 2^{21} clocks 10 2^{25} clocks 11 2^{29} clocks	
2:3	WRC	Watchdog Reset Control 00 No Watchdog reset will occur. 01 Core reset will be forced by the Watchdog. 10 Chip reset will be forced by the Watchdog. 11 System reset will be forced by the Watchdog.	TCR[WRC] resets to 00. This field can be set by software, but cannot be cleared by software, except by a software-induced reset.
4	WIE	Watchdog Interrupt Enable 0 Disable watchdog interrupt. 1 Enable watchdog interrupt.	
5	PIE	PIT Interrupt Enable 0 Disable PIT interrupt. 1 Enable PIT interrupt.	
6:7	FP	FIT Period 00 2^9 clocks 01 2^{13} clocks 10 2^{17} clocks 11 2^{21} clocks	
8	FIE	FIT Interrupt Enable 0 Disable FIT interrupt. 1 Enable FIT interrupt.	
9	ARE	Auto Reload Enable 0 Disable auto reload. 1 Enable auto reload.	Disables on reset.
10:31		Reserved	

SPR 0x3D8 Read/Clear

See “Timer Status Register (TSR)” on page 6-8.




Figure 10-34. Timer Status Register (TSR)

0	ENW	Enable Next Watchdog 0 Action on next watchdog event is to set TSR[ENW] = 1. 1 Action on next watchdog event is governed by TSR[WIS].	Software must reset TSR[ENW] = 0 after each watchdog timer event.
1	WIS	Watchdog Interrupt Status 0 No Watchdog interrupt is pending. 1 Watchdog interrupt is pending.	
2:3	WRS	Watchdog Reset Status 00 No Watchdog reset has occurred. 01 Core reset was forced by the watchdog. 10 Chip reset was forced by the watchdog. 11 System reset was forced by the watchdog.	
4	PIS	PIT Interrupt Status 0 No PIT interrupt is pending. 1 PIT interrupt is pending.	
5	FIS	FIT Interrupt Status 0 No FIT interrupt is pending. 1 FIT interrupt is pending.	
6:31		Reserved	

USPRG0

User Special Purpose Register General 0

SPR 0x100 (User R/W)

See “Special Purpose Register General (SPRG0–SPRG7)” on page 2-9.

0	31
---	----

Figure 10-35. User SPR General 0 (USPRG0)

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

SPR 0x001

See “Fixed Point Exception Register (XER)” on page 2-7.




Figure 10-36. Fixed Point Exception Register (XER)

0	SO	Summary Overflow 0 No overflow has occurred. 1 Overflow has occurred.	Can be <i>set</i> by mtspr or by using “o” form instructions; can be <i>reset</i> by mtspr or by mcrxr .
1	OV	Overflow 0 No overflow has occurred. 0 Overflow has occurred.	Can be <i>set</i> by mtspr or by using “o” form instructions; can be <i>reset</i> by mtspr , by mcrxr , or “o” form instructions.
2	CA	Carry 0 Carry has not occurred. 1 Carry has occurred.	Can be <i>set</i> by mtspr or arithmetic instructions that update the CA field; can be <i>reset</i> by mtspr , by mcrxr , or by arithmetic instructions that update the CA field.
3:24		Reserved	
25:31	TBC	Transfer Byte Count	Used by lswx and stswx ; written by mtspr .

ZPR

Zone Protection Register

SPR 0x3B0

See “Zone Protection” on page 7-14.




Figure 10-37. Zone Protection Register (ZPR)

0:1	Z0	TLB page access control for all pages in this zone. In the problem state (MSR[PR] = 1): 00 No access 01 Access controlled by applicable TLB_entry[EX, WR] 10 Access controlled by applicable TLB_entry[EX, WR] 11 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted In the supervisor state (MSR[PR] = 0): 00 Access controlled by applicable TLB_entry[EX, WR] 01 Access controlled by applicable TLB_entry[EX, WR] 10 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted 11 Accessed as if execute and write permissions (TLB_entry[EX, WR]) are granted
2:3	Z1	See the description of Z0.
4:5	Z2	See the description of Z0.
6:7	Z3	See the description of Z0.
8:9	Z4	See the description of Z0.
10:11	Z5	See the description of Z0.
12:13	Z6	See the description of Z0.
14:15	Z7	See the description of Z0.
16:17	Z8	See the description of Z0.
18:19	Z9	See the description of Z0.
20:21	Z10	See the description of Z0.
22:23	Z11	See the description of Z0.
24:25	Z12	See the description of Z0.
26:27	Z13	See the description of Z0.
28:29	Z14	See the description of Z0.
30:31	Z15	See the description of Z0.

Appendix A. Instruction Summary

This appendix contains PPC405 instructions summarized alphabetically and by opcode.

“Instruction Set and Extended Mnemonics – Alphabetical” lists all PPC405 mnemonics, including extended mnemonics, alphabetically. A short functional description is included for each mnemonic.

“Instructions Sorted by Opcode,” on page A-33, lists all PPC405 instructions, sorted by primary and secondary opcodes. Extended mnemonics are not included in the opcode list.

“Instruction Formats,” on page A-41, illustrates the PPC405 instruction forms (allowed arrangements of fields within instructions).

A.1 Instruction Set and Extended Mnemonics – Alphabetical

Table A-1 summarizes the PPC405 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. Chapter 9 also describes the instruction operands and notation.

Note the following for the branch conditional mnemonic:

Bit 4 of the BO field provides a hint about the most likely outcome of a conditional branch. (See “Branch Prediction” on page 2-26 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 the table below, $BO_4 = 0$ has always been assumed. The assembler must allow the programmer to specify branch prediction. To do this, the assembler supports a suffixes for the conditional branch mnemonics:

- + Predict branch to be taken.
- Predict branch not to be taken.

As specific examples, **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 2-26 for more information.

Table A-1. PPC405 Instruction Syntax Summary

Mnemonic	Operands	Function	Other Registers Changed	Page
add	RT, RA, RB	Add (RA) to (RB). Place result in RT.		9-6
add.			CR[CR0]	
addo			XER[SO, OV]	
addo.			CR[CR0] XER[SO, OV]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
addc	RT, RA, RB	Add (RA) to (RB). Place result in RT. Place carry-out in XER[CA].		9-7
addc.			CR[CR0]	
addco			XER[SO, OV]	
addco.			CR[CR0] XER[SO, OV]	
adde	RT, RA, RB	Add XER[CA], (RA), (RB). Place result in RT. Place carry-out in XER[CA].		9-8
adde.			CR[CR0]	
addeo			XER[SO, OV]	
addeo.			CR[CR0] XER[SO, OV]	
addi	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT.		9-9
addic	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].		9-10
addic.	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].	CR[CR0]	9-11
addis	RT, RA, IM	Add (IM ¹⁶ 0) to (RA 0). Place result in RT.		9-12
addme	RT, RA	Add XER[CA], (RA), (-1). Place result in RT. Place carry-out in XER[CA].		9-13
addme.			CR[CR0]	
addmeo			XER[SO, OV]	
addmeo.			CR[CR0] XER[SO, OV]	
addze	RT, RA	Add XER[CA] to (RA). Place result in RT. Place carry-out in XER[CA].		9-14
addze.			CR[CR0]	
addzeo			XER[SO, OV]	
addzeo.			CR[CR0] XER[SO, OV]	
and	RA, RS, RB	AND (RS) with (RB). Place result in RA.		9-15
and.			CR[CR0]	
andc	RA, RS, RB	AND (RS) with ¬(RB). Place result in RA.		9-16
andc.			CR[CR0]	
andi.	RA, RS, IM	AND (RS) with (¹⁶ 0 IM). Place result in RA.	CR[CR0]	9-17
andis.	RA, RS, IM	AND (RS) with (IM ¹⁶ 0). Place result in RA.	CR[CR0]	9-18

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
b	target	Branch unconditional relative. $LI \leftarrow (target - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^20)$		9-19
ba		Branch unconditional absolute. $LI \leftarrow target_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^20)$		
bl		Branch unconditional relative. $LI \leftarrow (target - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^20)$	$(LR) \leftarrow CIA + 4.$	
bla		Branch unconditional absolute. $LI \leftarrow target_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^20)$	$(LR) \leftarrow CIA + 4.$	
bc	BO, BI, target	Branch conditional relative. $BD \leftarrow (target - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$	9-20
bca		Branch conditional absolute. $BD \leftarrow target_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$	
bcl		Branch conditional relative. $BD \leftarrow (target - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$ $(LR) \leftarrow CIA + 4.$	
bcla		Branch conditional absolute. $BD \leftarrow target_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$ $(LR) \leftarrow CIA + 4.$	
bcctr	BO, BI	Branch conditional to address in CTR. Using (CTR) at exit from instruction, $NIA \leftarrow CTR_{0:29} \parallel ^20.$	CTR if $BO_2 = 0.$	9-26
bcctrl			CTR if $BO_2 = 0.$ $(LR) \leftarrow CIA + 4.$	
bclr	BO, BI	Branch conditional to address in LR. Using (LR) at entry to instruction, $NIA \leftarrow LR_{0:29} \parallel ^20.$	CTR if $BO_2 = 0.$	9-30
bctrl			CTR if $BO_2 = 0.$ $(LR) \leftarrow CIA + 4.$	
bctr		Branch unconditionally to address in CTR. <i>Extended mnemonic for</i> bctr 20,0		9-26
bctrl		<i>Extended mnemonic for</i> bctrl 20,0	$(LR) \leftarrow CIA + 4.$	
bdnz	target	Decrement CTR. Branch if $CTR \neq 0.$ <i>Extended mnemonic for</i> bc 16,0,target		9-20
bdnza		<i>Extended mnemonic for</i> bca 16,0,target		
bdnzl		<i>Extended mnemonic for</i> bcl 16,0,target	$(LR) \leftarrow CIA + 4.$	
bdnzla		<i>Extended mnemonic for</i> bcla 16,0,target	$(LR) \leftarrow CIA + 4.$	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bdnzlr		Decrement CTR. Branch if CTR \neq 0 to address in LR. <i>Extended mnemonic for</i> bclr 16,0		9-30
bdnzlrl		<i>Extended mnemonic for</i> bclrl 16,0	(LR) \leftarrow CIA + 4.	
bdnzf	cr_bit, target	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 0,cr_bit,target		9-20
bdnzfa		<i>Extended mnemonic for</i> bca 0,cr_bit,target		
bdnzfl		<i>Extended mnemonic for</i> bcl 0,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnzfla		<i>Extended mnemonic for</i> bcla 0,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnzflr	cr_bit	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 0,cr_bit		9-30
bdnzflrl		<i>Extended mnemonic for</i> bclrl 0,cr_bit	(LR) \leftarrow CIA + 4.	
bdnzf	cr_bit, target	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 8,cr_bit,target		9-20
bdnzta		<i>Extended mnemonic for</i> bca 8,cr_bit,target		
bdnztl		<i>Extended mnemonic for</i> bcl 8,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnztla		<i>Extended mnemonic for</i> bcla 8,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnztlr	cr_bit	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 1 to address in LR. <i>Extended mnemonic for</i> bclr 8,cr_bit		9-30
bdnztlrl		<i>Extended mnemonic for</i> bclrl 8,cr_bit	(LR) \leftarrow CIA + 4.	
bdz	target	Decrement CTR. Branch if CTR = 0. <i>Extended mnemonic for</i> bc 18,0,target		9-20
bdza		<i>Extended mnemonic for</i> bca 18,0,target		
bdzl		<i>Extended mnemonic for</i> bcl 18,0,target	(LR) \leftarrow CIA + 4.	
bdzla		<i>Extended mnemonic for</i> bcla 18,0,target	(LR) \leftarrow CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bdzlr		Decrement CTR. Branch if CTR = 0 to address in LR. <i>Extended mnemonic for</i> bclr 18,0		9-30
bdzlrI		<i>Extended mnemonic for</i> bclrI 18,0	(LR) ← CIA + 4.	
bdzfb	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 2,cr_bit,target		9-20
bdzfa		<i>Extended mnemonic for</i> bca 2,cr_bit,target		
bdzfbI		<i>Extended mnemonic for</i> bcl 2,cr_bit,target	(LR) ← CIA + 4.	
bdzfaI		<i>Extended mnemonic for</i> bcla 2,cr_bit,target	(LR) ← CIA + 4.	
bdzflr	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 2,cr_bit		9-30
bdzflrI		<i>Extended mnemonic for</i> bclrI 2,cr_bit	(LR) ← CIA + 4.	
bdztd	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 10,cr_bit,target		9-20
bdztda		<i>Extended mnemonic for</i> bca 10,cr_bit,target		
bdztdI		<i>Extended mnemonic for</i> bcl 10,cr_bit,target	(LR) ← CIA + 4.	
bdztdaI		<i>Extended mnemonic for</i> bcla 10,cr_bit,target	(LR) ← CIA + 4.	
bdztlr	cr_bit	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 1 to address in LR. <i>Extended mnemonic for</i> bclr 10,cr_bit		9-30
bdztlrI		<i>Extended mnemonic for</i> bclrI 10,cr_bit	(LR) ← CIA + 4.	
beq	[cr_field], target	Branch if equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+2,target		9-20
beqa		<i>Extended mnemonic for</i> bca 12,4*cr_field+2,target		
beqI		<i>Extended mnemonic for</i> bcl 12,4*cr_field+2,target	(LR) ← CIA + 4.	
beqaI		<i>Extended mnemonic for</i> bcla 12,4*cr_field+2,target	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
beqctr	[cr_field]	Branch if equal to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+2		9-26
beqctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+2	(LR) ← CIA + 4.	
beqlr	[cr_field]	Branch if equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+2		9-30
beqlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+2	(LR) ← CIA + 4.	
bf	cr_bit, target	Branch if CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 4,cr_bit,target		9-20
bfa		<i>Extended mnemonic for</i> bca 4,cr_bit,target		
bfl		<i>Extended mnemonic for</i> bcl 4,cr_bit,target	(LR) ← CIA + 4.	
bfla		<i>Extended mnemonic for</i> bcll 4,cr_bit,target	(LR) ← CIA + 4.	
bfctr	cr_bit	Branch if CR _{cr_bit} = 0 to address in CTR. <i>Extended mnemonic for</i> bcctr 4,cr_bit		9-26
bfctrl		<i>Extended mnemonic for</i> bcctrl 4,cr_bit	(LR) ← CIA + 4.	
bflr	cr_bit	Branch if CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 4,cr_bit		9-30
bflrl		<i>Extended mnemonic for</i> bclrl 4,cr_bit	(LR) ← CIA + 4.	
bge	[cr_field], target	Branch if greater than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target		9-20
bgea		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target		
bgel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bgela		<i>Extended mnemonic for</i> bcll 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bgectr	[cr_field]	Branch if greater than or equal to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0		9-26
bgectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bgelr	[cr_field]	Branch if greater than or equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0		9-30
bgelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bgt	[cr_field], target	Branch if greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+1,target		9-20
bgta		<i>Extended mnemonic for</i> bca 12,4*cr_field+1,target		
bgtl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+1,target	(LR) ← CIA + 4.	
bgtla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+1,target	(LR) ← CIA + 4.	
bgtctr	[cr_field]	Branch if greater than to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+1		9-26
bgtctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+1	(LR) ← CIA + 4.	
bgtlr	[cr_field]	Branch if greater than to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+1		9-30
bgtlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+1	(LR) ← CIA + 4.	
ble	[cr_field], target	Branch if less than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target		9-20
blea		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target		
blel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	(LR) ← CIA + 4.	
blela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	(LR) ← CIA + 4.	
blectr	[cr_field]	Branch if less than or equal to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1		9-26
blectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
blelr	[cr_field]	Branch if less than or equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1		9-30
blelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) ← CIA + 4.	
blr		Branch unconditionally to address in LR. <i>Extended mnemonic for</i> bclr 20,0		9-30
blrl		<i>Extended mnemonic for</i> bclrl 20,0	(LR) ← CIA + 4.	
blt	[cr_field], target	Branch if less than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+0,target		9-20
blta		<i>Extended mnemonic for</i> bca 12,4*cr_field+0,target		
bltl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+0,target	(LR) ← CIA + 4.	
bltla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+0,target	(LR) ← CIA + 4.	
blctr	[cr_field]	Branch if less than to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+0		9-26
blctrll		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+0	(LR) ← CIA + 4.	
bltlr	[cr_field]	Branch if less than to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+0		9-30
bltlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+0	(LR) ← CIA + 4.	
bne	[cr_field], target	Branch if not equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+2,target		9-20
bnea		<i>Extended mnemonic for</i> bca 4,4*cr_field+2,target		
bnel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+2,target	(LR) ← CIA + 4.	
bnela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+2,target	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bnctr	[cr_field]	Branch if not equal to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+2		9-26
bnctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+2	(LR) ← CIA + 4.	
bnelr	[cr_field]	Branch if not equal to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+2		9-30
bnelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+2	(LR) ← CIA + 4.	
bng	[cr_field], target	Branch if not greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target		9-20
bnga		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target		
bngl		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	(LR) ← CIA + 4.	
bngla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	(LR) ← CIA + 4.	
bngctr	[cr_field]	Branch if not greater than to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1		9-26
bngctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.	
bnglr	[cr_field]	Branch if not greater than to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1		9-30
bnglrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) ← CIA + 4.	
bnl	[cr_field], target	Branch if not less than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target		9-20
bnla		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target		
bnll		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bnlla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+0,target	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bnlctr	[cr_field]	Branch if not less than to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0		9-26
bnlctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bnllr	[cr_field]	Branch if not less than to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0		9-30
bnllrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bns	[cr_field], target	Branch if not summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target		9-20
bnsa		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target		
bnsi		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnsia		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnsctr	[cr_field]	Branch if not summary overflow to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3		9-26
bnsctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnslr	[cr_field]	Branch if not summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3		9-30
bnsrlr		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnu	[cr_field], target	Branch if not unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target		9-20
bnu		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target		
bnul		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnula		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bnuctr	[cr_field]	Branch if not unordered to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3		9-26
bnuctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnulr	[cr_field]	Branch if not unordered to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3		9-30
bnulrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bso	[cr_field], target	Branch if summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target		9-20
boas		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target		
bsol		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bsola		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bsoctr	[cr_field]	Branch if summary overflow to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3		9-26
bsoctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) ← CIA + 4.	
bsolr	[cr_field]	Branch if summary overflow to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3		9-30
bsolrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) ← CIA + 4.	
bt	cr_bit, target	Branch if CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 12,cr_bit,target		9-20
bta		<i>Extended mnemonic for</i> bca 12,cr_bit,target		
btl		<i>Extended mnemonic for</i> bcl 12,cr_bit,target	(LR) ← CIA + 4.	
btla		<i>Extended mnemonic for</i> bcla 12,cr_bit,target	(LR) ← CIA + 4.	
btctr	cr_bit	Branch if CR _{cr_bit} = 1 to address in CTR. <i>Extended mnemonic for</i> bcctr 12,cr_bit		9-26
btctrl		<i>Extended mnemonic for</i> bcctrl 12,cr_bit	(LR) ← CIA + 4.	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
btlr	cr_bit	Branch if $CR_{cr_bit} = 1$, to address in LR. <i>Extended mnemonic for</i> bclr 12,cr_bit		9-30
btlrl		<i>Extended mnemonic for</i> bclrl 12,cr_bit	(LR) \leftarrow CIA + 4.	
bun	[cr_field], target	Branch if unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target		9-20
buna		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target		
bunl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) \leftarrow CIA + 4.	
bunla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) \leftarrow CIA + 4.	
bunctr	[cr_field]	Branch if unordered to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3		9-26
bunctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) \leftarrow CIA + 4.	
bunlr	[cr_field]	Branch if unordered, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3		9-30
bunlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) \leftarrow CIA + 4.	
clrlwi	RA, RS, n	Clear left immediate. ($n < 32$) $(RA)_{0:n-1} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,n,31		9-147
clrlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,n,31	CR[CR0]	
clrlslwi	RA, RS, b, n	Clear left and shift left immediate. ($n \leq b < 32$) $(RA)_{b-n:31-n} \leftarrow (RS)_{b:31}$ $(RA)_{32-n:31} \leftarrow {}^n0$ $(RA)_{0:b-n-1} \leftarrow {}^{b-n}0$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,b-n,31-n		9-147
clrlslwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,b-n,31-n	CR[CR0]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
clrrwi	RA, RS, n	Clear right immediate. ($n < 32$) $(RA)_{32-n:31} \leftarrow n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,0,31-n		9-147
clrrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,0,31-n	CR[CR0]	
cmp	BF, 0, RA, RB	Compare (RA) to (RB), signed. Results in CR[CRn], where $n = BF$.		9-34
cmpi	BF, 0, RA, IM	Compare (RA) to EXTS(IM), signed. Results in CR[CRn], where $n = BF$.		9-35
cmpl	BF, 0, RA, RB	Compare (RA) to (RB), unsigned. Results in CR[CRn], where $n = BF$.		9-36
cmpli	BF, 0, RA, IM	Compare (RA) to ($^{16}0 \parallel IM$), unsigned. Results in CR[CRn], where $n = BF$.		9-37
cmplw	[BF,] RA, RB	Compare Logical Word. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmpl BF,0,RA,RB		9-36
cmplwi	[BF,] RA, IM	Compare Logical Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmpli BF,0,RA,IM		9-37
cmpw	[BF,] RA, RB	Compare Word. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmp BF,0,RA,RB		9-34
cmpwi	[BF,] RA, IM	Compare Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for</i> cmpi BF,0,RA,IM		9-35
cntlzw	RA, RS	Count leading zeros in RS. Place result in RA.		9-38
cntlzw.			CR[CR0]	
crand	BT, BA, BB	AND bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-39
crandc	BT, BA, BB	AND bit (CR_{BA}) with $\neg(CR_{BB})$. Place result in CR_{BT} .		9-40
crclr	bx	Condition register clear. <i>Extended mnemonic for</i> crxor bx,bx,bx		9-46
creqv	BT, BA, BB	Equivalence of bit CR_{BA} with CR_{BB} . $CR_{BT} \leftarrow \neg(CR_{BA} \oplus CR_{BB})$		9-41
crmove	bx, by	Condition register move. <i>Extended mnemonic for</i> cror bx,by,by		9-44
crnand	BT, BA, BB	NAND bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-42

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
cnor	BT, BA, BB	NOR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-43
crnot	bx, by	Condition register not. <i>Extended mnemonic for</i> cnor bx,by,by		9-43
cror	BT, BA, BB	OR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-44
crorc	BT, BA, BB	OR bit (CR_{BA}) with $\neg(CR_{BB})$. Place result in CR_{BT} .		9-45
crset	bx	Condition register set. <i>Extended mnemonic for</i> creqv bx,bx,bx		9-41
crxor	BT, BA, BB	XOR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-46
dcba	RA, RB	Speculatively establish the data cache block which contains the effective address $(RA 0) + (RB)$.		9-47
dcbf	RA, RB	Flush (store, then invalidate) the data cache block which contains the effective address $(RA 0) + (RB)$.		9-49
dcbi	RA, RB	Invalidate the data cache block which contains the effective address $(RA 0) + (RB)$.		9-50
dcbst	RA, RB	Store the data cache block which contains the effective address $(RA 0) + (RB)$.		9-51
dcbt	RA, RB	Load the data cache block which contains the effective address $(RA 0) + (RB)$.		9-52
dcbtst	RA, RB	Load the data cache block which contains the effective address $(RA 0) + (RB)$.		9-53
dcbz	RA, RB	Zero the data cache block which contains the effective address $(RA 0) + (RB)$.		9-54
dccci	RA, RB	Invalidate the data cache congruence class associated with the effective address $(RA 0) + (RB)$.		9-56
dcread	RT, RA, RB	Read either tag or data information from the data cache congruence class associated with the effective address $(RA 0) + (RB)$. Place the results in RT.		9-57
divw	RT, RA, RB	Divide (RA) by (RB), signed. Place result in RT.		9-59
divw.			CR[CR0]	
divwo			XER[SO, OV]	
divwo.			CR[CR0] XER[SO, OV]	
divwu	RT, RA, RB	Divide (RA) by (RB), unsigned. Place result in RT.		9-60
divwu.			CR[CR0]	
divwuo			XER[SO, OV]	
divwuo.			CR[CR0] XER[SO, OV]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
eieio		Storage synchronization. All loads and stores that precede the eieio instruction complete before any loads and stores that follow the instruction access main storage. Implemented as sync , which is more restrictive.		9-61
eqv	RA, RS, RB	Equivalence of (RS) with (RB). $(RA) \leftarrow \neg((RS) \oplus (RB))$		9-62
eqv.			CR[CR0]	
extlwi	RA, RS, n, b	Extract and left justify immediate. ($n > 0$) $(RA)_{0:n-1} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{n:31} \leftarrow 32-n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,b,0,n-1		9-60
extlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,b,0,n-1	CR[CR0]	
extrwi	RA, RS, n, b	Extract and right justify immediate. ($n > 0$) $(RA)_{32-n:31} \leftarrow (RS)_{b:b+n-1}$ $(RA)_{0:31-n} \leftarrow 32-n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,b+n,32-n,31		9-147
extrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,b+n,32-n,31	CR[CR0]	
extsb	RA, RS	Extend the sign of byte (RS) _{24:31} . Place the result in RA.		9-63
extsb.			CR[CR0]	
extsh	RA, RS	Extend the sign of halfword (RS) _{16:31} . Place the result in RA.		9-64
extsh.			CR[CR0]	
icbi	RA, RB	Invalidate the instruction cache block which contains the effective address (RA 0) + (RB).		9-65
icbt	RA, RB	Load the instruction cache block which contains the effective address (RA 0) + (RB).		9-66
iccci	RA, RB	Invalidate instruction cache.		9-67
icread	RA, RB	Read either tag or data information from the instruction cache congruence class associated with the effective address (RA 0) + (RB). Place the results in ICDBDR.		9-68
inslwi	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> rlwimi RA,RS,32-b,b,b+n-1		9-146
inslwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b,b,b+n-1	CR[CR0]	
insrwi	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> rlwimi RA,RS,32-b-n,b,b+n-1		9-146
insrwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b-n,b,b+n-1	CR[CR0]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
isync		Synchronize execution context by flushing the prefetch queue.		9-70
la	RT, D(RA)	Load address. ($RA \neq 0$) 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> addi RT,RA,D		9-9
lbz	RT, D(RA)	Load byte from $EA = (RA 0) + \text{EXTS}(D)$ and pad left with zeroes, $(RT) \leftarrow {}^{24}0 \parallel \text{MS}(EA,1)$.		9-71
lbzu	RT, D(RA)	Load byte from $EA = (RA 0) + \text{EXTS}(D)$ and pad left with zeroes, $(RT) \leftarrow {}^{24}0 \parallel \text{MS}(EA,1)$. Update the base address, $(RA) \leftarrow EA$.		9-72
lbzux	RT, RA, RB	Load byte from $EA = (RA 0) + (RB)$ and pad left with zeroes, $(RT) \leftarrow {}^{24}0 \parallel \text{MS}(EA,1)$. Update the base address, $(RA) \leftarrow EA$.		9-73
lbzx	RT, RA, RB	Load byte from $EA = (RA 0) + (RB)$ and pad left with zeroes, $(RT) \leftarrow {}^{24}0 \parallel \text{MS}(EA,1)$.		9-74
lha	RT, D(RA)	Load halfword from $EA = (RA 0) + \text{EXTS}(D)$ and sign extend, $(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$.		9-75
lhau	RT, D(RA)	Load halfword from $EA = (RA 0) + \text{EXTS}(D)$ and sign extend, $(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$. Update the base address, $(RA) \leftarrow EA$.		9-76
lhaux	RT, RA, RB	Load halfword from $EA = (RA 0) + (RB)$ and sign extend, $(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$. Update the base address, $(RA) \leftarrow EA$.		9-77
lhax	RT, RA, RB	Load halfword from $EA = (RA 0) + (RB)$ and sign extend, $(RT) \leftarrow \text{EXTS}(\text{MS}(EA,2))$.		9-78
lhbrx	RT, RA, RB	Load halfword from $EA = (RA 0) + (RB)$, then reverse byte order and pad left with zeroes, $(RT) \leftarrow {}^{16}0 \parallel \text{MS}(EA+1,1) \parallel \text{MS}(EA,1)$.		9-79
lhz	RT, D(RA)	Load halfword from $EA = (RA 0) + \text{EXTS}(D)$ and pad left with zeroes, $(RT) \leftarrow {}^{16}0 \parallel \text{MS}(EA,2)$.		9-80

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
lhz	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) \leftarrow $^{16}0 \parallel$ MS(EA,2). Update the base address, (RA) \leftarrow EA.		9-81
lhzx	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{16}0 \parallel$ MS(EA,2). Update the base address, (RA) \leftarrow EA.		9-82
lhzx	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{16}0 \parallel$ MS(EA,2).		9-83
li	RT, IM	Load immediate. (RT) \leftarrow EXTS(IM) <i>Extended mnemonic for</i> addi RT,0,value		9-9
lis	RT, IM	Load immediate shifted. (RT) \leftarrow (IM \parallel $^{16}0$) <i>Extended mnemonic for</i> addis RT,0,value		9-12
lmw	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).		9-84
lswi	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 _{FINAL} \leftarrow ((RT + CEIL(n/4) - 1) % 32). GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R _{FINAL} .		9-85
lswx	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 _{FINAL} \leftarrow ((RT + CEIL(n/4) - 1) % 32). GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R _{FINAL} . RB is not altered unless RB = R _{FINAL} . If n=0, content of RT is undefined.		9-87
lwarx	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) \leftarrow MS(EA,4). Set the Reservation bit.		9-89
lwbrx	RT, RA, RB	Load word from EA = (RA 0) + (RB) then reverse byte order, (RT) \leftarrow MS(EA+3,1) \parallel MS(EA+2,1) \parallel MS(EA+1,1) \parallel MS(EA,1).		9-90

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
lwz	RT, D(RA)	Load word from EA = (RA 0) + EXTS(D) and place in RT, (RT) ← MS(EA,4).		9-91
lwzu	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.		9-92
lwzux	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4). Update the base address, (RA) ← EA.		9-93
lwzx	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4).		9-94
macchw	RT, RA, RB	prod _{0:31} ← (RA) _{16:31} × (RB) _{0:15} signed temp _{0:32} ← prod _{0:31} + (RT) (RT) ← temp _{1:32}		9-95
macchw.			CR[CR0]	
macchwo			XER[SO, OV]	
macchwo.			CR[CR0] XER[SO, OV]	
macchws	RT, RA, RB	prod _{0:31} ← (RA) _{16:31} × (RB) _{0:15} signed temp _{0:32} ← prod _{0:31} + (RT) if ((prod ₀ = RT ₀) ∧ (RT ₀ ≠ temp ₁)) then (RT) ← (RT ₀ ³¹ (¬RT ₀)) else (RT) ← temp _{1:32}		9-96
macchws.			CR[CR0]	
macchwso			XER[SO, OV]	
macchwso.			CR[CR0] XER[SO, OV]	
macchwsu	RT, RA, RB	prod _{0:31} ← (RA) _{16:31} × (RB) _{0:15} unsigned temp _{0:32} ← prod _{0:31} + (RT) (RT) ← (temp _{1:32} ∨ ³² temp ₀)		9-97
macchwsu.			CR[CR0]	
macchwsuo			XER[SO, OV]	
macchwsuo.			CR[CR0] XER[SO, OV]	
macchwu	RT, RA, RB	prod _{0:31} ← (RA) _{16:31} × (RB) _{0:15} unsigned temp _{0:32} ← prod _{0:31} + (RT) (RT) ← temp _{1:32}		9-98
macchwu.			CR[CR0]	
macchwuo			XER[SO, OV]	
macchwuo.			CR[CR0] XER[SO, OV]	
machhw	RT, RA, RB	prod _{0:15} ← (RA) _{16:31} × (RB) _{0:15} signed temp _{0:32} ← prod _{0:31} + (RT) (RT) ← temp _{1:32}		9-99
machhw.			CR[CR0]	
machhwo			XER[SO, OV]	
machhwo.			CR[CR0] XER[SO, OV]	
machhws	RT, RA, RB	prod _{0:31} ← (RA) _{0:15} × (RB) _{0:15} signed temp _{0:32} ← prod _{0:31} + (RT) if ((prod ₀ = RT ₀) ∧ (RT ₀ ≠ temp ₁)) then (RT) ← (RT ₀ ³¹ (¬RT ₀)) else (RT) ← temp _{1:32}		9-100
machhws.			CR[CR0]	
machhwso			XER[SO, OV]	
machhwso.			CR[CR0] XER[SO, OV]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
machwsu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee {}^{32}temp_0)$		9-101
machwsu.			CR[CR0]	
machwsuo			XER[SO, OV]	
machwsuo.			CR[CR0] XER[SO, OV]	
machwu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-102
machwu.			CR[CR0]	
machwuo			XER[SO, OV]	
machwuo.			CR[CR0] XER[SO, OV]	
maclhw	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ signed $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-103
maclhw.			CR[CR0]	
maclhwo			XER[SO, OV]	
maclhwo.			CR[CR0] XER[SO, OV]	
maclhws	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ signed $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ if $((prod_0 = RT_0) \wedge (RT_0 \neq temp_1))$ then $(RT) \leftarrow (RT_0 \parallel {}^{31}(-RT_0))$ else $(RT) \leftarrow temp_{1:32}$		9-104
maclhws.			CR[CR0]	
maclhwso			XER[SO, OV]	
maclhwso.			CR[CR0] XER[SO, OV]	
maclhwsu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee {}^{32}temp_0)$		9-105
maclhwsu.			CR[CR0]	
maclhwsuo			XER[SO, OV]	
maclhwsuo.			CR[CR0] XER[SO, OV]	
maclhwu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-106
maclhwu.			CR[CR0]	
maclhwuo			XER[SO, OV]	
maclhwuo.			CR[CR0] XER[SO, OV]	
mcrf	BF, BFA	Move CR field, $(CR[CR_n]) \leftarrow (CR[CR_m])$ where $m \leftarrow BFA$ and $n \leftarrow BF$.		9-107
mcrxr	BF	Move XER[0:3] into field CR _n , where $n \leftarrow BF$. $CR[CR_n] \leftarrow (XER[SO, OV, CA])$. $(XER[SO, OV, CA]) \leftarrow {}^30$.		9-108
mfcr	RT	Move from CR to RT, $(RT) \leftarrow (CR)$.		9-109
mfocr	RT, DCRN	Move from DCR to RT, $(RT) \leftarrow (DCR(DCRN))$.		9-110
mfmsr	RT	Move from MSR to RT, $(RT) \leftarrow (MSR)$.		9-111

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mfccr0 mfctr mfdac1 mfdac2 mfdear mfdbcr0 mfdbcr1 mfdbsr mfdccr mfdcwr mfdcv1 mfdcv2 mfesr mfevpr mfiac1 mfiac2 mfiac3 mfiac4 mficcr mficbdr mflr mfpid mfpit mfpvr mfsgl mfslr mfsprg0 mfsprg1 mfsprg2 mfsprg3 mfsprg4 mfsprg5 mfsprg6 mfsprg7 mfsrc0 mfsrc1 mfsrc2 mfsrc3 mfsrc0r mftcr mftsr mfscr mfzpr	RT	Move from special purpose register (SPR) SPRN. <i>Extended mnemonic for</i> mfscr RT,SPRN See Table 10.5, "Special Purpose Registers," on page 10-2 for listing of valid SPRN values.		9-112
mfscr	RT, SPRN	Move from SPR to RT, (RT) ← (SPR(SPRN)).		9-112
mftb	RT, TBRN	Move from TBR to RT, (RT) ← (TBR(TBRN)).		9-114
mftb	RT	Move the contents of TBL into RT, (RT) ← (TBL) <i>Extended mnemonic for</i> mftb RT,TBL		9-114

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mftbu	RT	Move the contents of TBU into RT, (RT) ← (TBU) <i>Extended mnemonic for mftb RT,TBU</i>		9-114
mr	RT, RS	Move register. (RT) ← (RS) <i>Extended mnemonic for or RT,RS,RS</i>		9-140
mr.		<i>Extended mnemonic for or. RT,RS,RS</i>	CR[CR0]	
mtcr	RS	Move to Condition Register. <i>Extended mnemonic for mtcrf 0xFF,RS</i>		9-116
mtcrf	FXM, RS	Move some or all of the contents of RS into CR as specified by FXM field, mask ← ${}^4(\text{FXM}_0) \parallel {}^4(\text{FXM}_1) \parallel \dots \parallel$ ${}^4(\text{FXM}_6) \parallel {}^4(\text{FXM}_7)$. (CR) ← ((RS) ∧ mask) ∨ (CR) ∧ ¬mask).		9-116
mtdcr	DCRN, RS	Move to DCR from RS, (DCR(DCRN)) ← (RS).		9-117
mtmsr	RS	Move to MSR from RS, (MSR) ← (RS).		9-118

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mtccr0 mtctr mtdac1 mtdac2 mtdbcr0 mtdbcr1 mtdbsr mtdccr mtdear mtdcwr mtdvc1 mtdvc2 mtesr mtevpr mtiac1 mtiac2 mtiac3 mtiac4 mticcr mticdbdr mtlr mtpid mtpit mtpvr mtsgr mtsler mtsprg0 mtsprg1 mtsprg2 mtsprg3 mtsprg4 mtsprg5 mtsprg6 mtsprg7 mtsrr0 mtsrr1 mtsrr2 mtsrr3 mtsu0r mttbl mttbu mttcr mttsr mtxer mtzpr	RS	Move to SPR SPRN. <i>Extended mnemonic for mtspr SPRN,RS</i> See Table 10.5, "Special Purpose Registers," on page 10-2 for listing of valid SPRN values.		9-119
mtspr	SPRN, RS	Move to SPR from RS, (SPR(SPRN)) ← (RS).		9-119
mulchw	RT, RA, RB	(RT) _{0:31} ← (RA) _{16:31} × (RB) _{0:15} signed		9-121
mulchw.			CR[CR0]	
mulchwu	RT, RA, RB	(RT) _{0:31} ← (RA) _{16:31} × (RB) _{0:15} unsigned		9-122
mulchwu.			CR[CR0]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mulhhw	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ signed		9-123
mulhhw.			CR[CR0]	
mulhhwu	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned		9-124
mulhhwu.			CR[CR0]	
mullhw	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ signed		9-125
mullhw.			CR[CR0]	
mullhwu	RT, RA, RB	$(RT)_{16:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned		9-126
mullhwu.			CR[CR0]	
mulhw	RT, RA, RB	Multiply (RA) and (RB), signed. Place high-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). $(RT) \leftarrow prod_{0:31}$.		9-127
mulhw.			CR[CR0]	
mulhwu	RT, RA, RB	Multiply (RA) and (RB), unsigned. Place high-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (unsigned). $(RT) \leftarrow prod_{0:31}$.		9-128
mulhwu.			CR[CR0]	
mulli	RT, RA, IM	Multiply (RA) and IM, signed. Place low-order result in RT. $prod_{0:47} \leftarrow (RA) \times IM$ (signed) $(RT) \leftarrow prod_{16:47}$		9-129
mullw	RT, RA, RB	Multiply (RA) and (RB), signed. Place low-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). $(RT) \leftarrow prod_{32:63}$.		9-130
mullw.			CR[CR0]	
mullwo			XER[SO, OV]	
mullwo.			CR[CR0] XER[SO, OV]	
nand	RA, RS, RB	NAND (RS) with (RB). Place result in RA.		9-131
nand.			CR[CR0]	
neg	RT, RA	Negative (twos complement) of RA. $(RT) \leftarrow \neg(RA) + 1$		9-132
neg.			CR[CR0]	
nego			XER[SO, OV]	
nego.			CR[CR0] XER[SO, OV]	
nmacchw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15})$ signed $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-133
nmacchw.			CR[CR0]	
nmacchwso			XER[SO, OV]	
nmacchwso.			CR[CR0] XER[SO, OV]	
nmacchws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15})$ 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 $(RT) \leftarrow temp_{1:32}$		9-134
nmacchws.			CR[CR0]	
nmacchwso			XER[SO, OV]	
nmacchwso.			CR[CR0] XER[SO, OV]	

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
nmachhw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ signed $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-135
nmachhw.			CR[CR0]	
nmachwso			XER[SO, OV]	
nmachwso.			CR[CR0] XER[SO, OV]	
nmachhws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ 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 $(RT) \leftarrow temp_{1:32}$		9-137
nmachhws.			CR[CR0]	
nmachhwso			XER[SO, OV]	
nmachhwso.			CR[CR0] XER[SO, OV]	
nmachlw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{16:31})$ 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 $(RT) \leftarrow temp_{1:32}$		9-138
nmachlw.			CR[CR0]	
nmachlwso			XER[SO, OV]	
nmachlwso.			CR[CR0] XER[SO, OV]	
nmachlws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ 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 $(RT) \leftarrow temp_{1:32}$		9-136
nmachlws.			CR[CR0]	
nmachlwso			XER[SO, OV]	
nmachlwso.			CR[CR0] XER[SO, OV]	
nop		Preferred no-op, triggers optimizations based on no-ops. <i>Extended mnemonic for ori 0,0,0</i>		9-134
nor	RA, RS, RB	NOR (RS) with (RB). Place result in RA.		9-139
nor.			CR[CR0]	
not	RA, RS	Complement register. $(RA) \leftarrow \neg(RS)$ <i>Extended mnemonic for nor RA,RS,RS</i>		9-139
not.			<i>Extended mnemonic for nor. RA,RS,RS</i> CR[CR0]	
or	RA, RS, RB	OR (RS) with (RB). Place result in RA.		9-134
or.			CR[CR0]	
orc	RA, RS, RB	OR (RS) with $\neg(RB)$. Place result in RA.		9-134
orc.			CR[CR0]	
ori	RA, RS, IM	OR (RS) with $(^{16}0 \parallel IM)$. Place result in RA.		9-134
oris	RA, RS, IM	OR (RS) with $(IM \parallel ^{16}0)$. Place result in RA.		9-143
rfei		Return from critical interrupt $(PC) \leftarrow (SRR2)$. $(MSR) \leftarrow (SRR3)$.		9-144

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
rfi		Return from interrupt. (PC) \leftarrow (SRR0). (MSR) \leftarrow (SRR1).		9-145
rlwimi rlwimi.	RA, RS, SH, MB, ME	Rotate left word immediate, then insert according to mask. $r \leftarrow \text{ROTL}((RS), SH)$ $m \leftarrow \text{MASK}(MB, ME)$ (RA) $\leftarrow (r \wedge m) \vee ((RA) \wedge \neg m)$	CR[CR0]	9-146
rlwinm rlwinm.	RA, RS, SH, MB, ME	Rotate left word immediate, then AND with mask. $r \leftarrow \text{ROTL}((RS), SH)$ $m \leftarrow \text{MASK}(MB, ME)$ (RA) $\leftarrow (r \wedge m)$	CR[CR0]	9-147
rlwnm rlwnm.	RA, RS, RB, MB, ME	Rotate left word, then AND with mask. $r \leftarrow \text{ROTL}((RS), (RB)_{27:31})$ $m \leftarrow \text{MASK}(MB, ME)$ (RA) $\leftarrow (r \wedge m)$	CR[CR0]	9-150
rotlw	RA, RS, RB	Rotate left. (RA) $\leftarrow \text{ROTL}((RS), (RB)_{27:31})$ <i>Extended mnemonic for</i> rlwnm RA,RS,RB,0,31		9-150
rotlw.		<i>Extended mnemonic for</i> rlwnm. RA,RS,RB,0,31	CR[CR0]	
rotlwi	RA, RS, n	Rotate left immediate. (RA) $\leftarrow \text{ROTL}((RS), n)$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,0,31		9-147
rotlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31	CR[CR0]	
rotrwi	RA, RS, n	Rotate right immediate. (RA) $\leftarrow \text{ROTL}((RS), 32-n)$ <i>Extended mnemonic for</i> rlwinm RA,RS,32-n,0,31		9-147
rotrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,0,31	CR[CR0]	
sc		System call exception is generated. (SRR1) \leftarrow (MSR) (SRR0) \leftarrow (PC) PC \leftarrow EVPR _{0:15} x'0C00' (MSR[WE, PR, EE, PE, DR, IR]) \leftarrow 0		9-151
slw slw.	RA, RS, RB	Shift left (RS) by (RB) _{27:31} . $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), n)$. if (RB) ₂₆ = 0 then $m \leftarrow \text{MASK}(0, 31 - n)$ else $m \leftarrow 32_0$. (RA) $\leftarrow r \wedge m$.	CR[CR0]	9-152

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
slwi	RA, RS, n	Shift left immediate. ($n < 32$) $(RA)_{0:31-n} \leftarrow (RS)_{n:31}$ $(RA)_{32-n:31} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,0,31-n		9-147
slwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31-n	CR[CR0]	
sraw	RA, RS, RB	Shift right algebraic (RS) by $(RB)_{27:31}$. $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. if $(RB)_{26} = 0$ then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow {}^{32}0$. $s \leftarrow (RS)_0$. $(RA) \leftarrow (r \wedge m) \vee ({}^{32}s \wedge \neg m)$. $\text{XER[CA]} \leftarrow s \wedge ((r \wedge \neg m) \neq 0)$.	CR[CR0]	9-153
sraw.				
srawi	RA, RS, SH	Shift right algebraic (RS) by SH. $n \leftarrow \text{SH}$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. $m \leftarrow \text{MASK}(n, 31)$. $s \leftarrow (RS)_0$. $(RA) \leftarrow (r \wedge m) \vee ({}^{32}s \wedge \neg m)$. $\text{XER[CA]} \leftarrow s \wedge ((r \wedge \neg m) \neq 0)$.	CR[CR0]	9-154
srawi.				
srw	RA, RS, RB	Shift right (RS) by $(RB)_{27:31}$. $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. if $(RB)_{26} = 0$ then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow {}^{32}0$. $(RA) \leftarrow r \wedge m$.	CR[CR0]	9-155
srw.				
srwi	RA, RS, n	Shift right immediate. ($n < 32$) $(RA)_{n:31} \leftarrow (RS)_{0:31-n}$ $(RA)_{0:n-1} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,32-n,n,31		9-147
srwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,n,31	CR[CR0]	
stb	RS, D(RA)	Store byte $(RS)_{24:31}$ in memory at $EA = (RA 0) + \text{EXTS}(D)$.		9-156
stbu	RS, D(RA)	Store byte $(RS)_{24:31}$ in memory at $EA = (RA 0) + \text{EXTS}(D)$. Update the base address, $(RA) \leftarrow EA$.		9-157
stbux	RS, RA, RB	Store byte $(RS)_{24:31}$ in memory at $EA = (RA 0) + (RB)$. Update the base address, $(RA) \leftarrow EA$.		9-158
stbx	RS, RA, RB	Store byte $(RS)_{24:31}$ in memory at $EA = (RA 0) + (RB)$.		9-159
sth	RS, D(RA)	Store halfword $(RS)_{16:31}$ in memory at $EA = (RA 0) + \text{EXTS}(D)$.		9-160

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
sthbrx	RS, RA, RB	Store halfword (RS) _{16:31} byte-reversed in memory at EA = (RA 0) + (RB). MS(EA, 2) ← (RS) _{24:31} (RS) _{16:23}		9-161
sthu	RS, D(RA)	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + EXT(S,D). Update the base address, (RA) ← EA.		9-162
sthux	RS, RA, RB	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		9-163
sthx	RS, RA, RB	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + (RB).		9-164
stmw	RS, D(RA)	Store consecutive words from RS through GPR(31) in memory starting at EA = (RA 0) + EXT(S,D).		9-165
stswi	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).		9-166
stswx	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).		9-167
stw	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXT(S,D).		9-169
stwbrx	RS, RA, RB	Store word (RS) byte-reversed in memory at EA = (RA 0) + (RB). MS(EA, 4) ← (RS) _{24:31} (RS) _{16:23} (RS) _{8:15} (RS) _{0:7}		9-170
stwcx.	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]) ← ² 0 1 XER _{so} else (CR[CR0]) ← ² 0 0 XER _{so} .		9-171
stwu	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXT(S,D). Update the base address, (RA) ← EA.		9-173

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
stwux	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		9-174
stwx	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB).		9-175
sub	RT, RA, RB	Subtract (RB) from (RA). (RT) ← ¬(RB) + (RA) + 1. <i>Extended mnemonic for subf RT,RB,RA</i>		9-176
sub.		<i>Extended mnemonic for subf. RT,RB,RA</i>	CR[CR0]	
subo		<i>Extended mnemonic for subfo RT,RB,RA</i>	XER[SO, OV]	
subo.		<i>Extended mnemonic for subfo. RT,RB,RA</i>	CR[CR0] XER[SO, OV]	
subc	RT, RA, RB	Subtract (RB) from (RA). (RT) ← ¬(RB) + (RA) + 1. Place carry-out in XER[CA]. <i>Extended mnemonic for subfc RT,RB,RA</i>		9-177
subc.		<i>Extended mnemonic for subfc. RT,RB,RA</i>	CR[CR0]	
subco		<i>Extended mnemonic for subfco RT,RB,RA</i>	XER[SO, OV]	
subco.		<i>Extended mnemonic for subfco. RT,RB,RA</i>	CR[CR0] XER[SO, OV]	
subf	RT, RA, RB	Subtract (RA) from (RB). (RT) ← ¬(RA) + (RB) + 1.		9-176
subf.			CR[CR0]	
subfo			XER[SO, OV]	
subfo.			CR[CR0] XER[SO, OV]	
subfc	RT, RA, RB	Subtract (RA) from (RB). (RT) ← ¬(RA) + (RB) + 1. Place carry-out in XER[CA].		9-177
subfc.			CR[CR0]	
subfco			XER[SO, OV]	
subfco.			CR[CR0] XER[SO, OV]	
subfe	RT, RA, RB	Subtract (RA) from (RB) with carry-in. (RT) ← ¬(RA) + (RB) + XER[CA]. Place carry-out in XER[CA].		9-178
subfe.			CR[CR0]	
subfeo			XER[SO, OV]	
subfeo.			CR[CR0] XER[SO, OV]	
subfic	RT, RA, IM	Subtract (RA) from EXTS(IM). (RT) ← ¬(RA) + EXTS(IM) + 1. Place carry-out in XER[CA].		9-179

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
subfme	RT, RA, RB	Subtract (RA) from (-1) with carry-in. (RT) $\leftarrow \neg(RA) + (-1) + XER[CA]$. Place carry-out in XER[CA].		9-180
subfme.			CR[CR0]	
subfmeo			XER[SO, OV]	
subfmeo.			CR[CR0] XER[SO, OV]	
subfze	RT, RA, RB	Subtract (RA) from zero with carry-in. (RT) $\leftarrow \neg(RA) + XER[CA]$. Place carry-out in XER[CA].		9-181
subfze.			CR[CR0]	
subfzeo			XER[SO, OV]	
subfzeo.			CR[CR0] XER[SO, OV]	
subi	RT, RA, IM	Subtract EXTS(IM) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> addi RT,RA,-IM		9-9
subic	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> addic RT,RA,-IM		9-10
subic.	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> addic. RT,RA,-IM	CR[CR0]	9-11
subis	RT, RA, IM	Subtract (IM ¹⁶ 0) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> addis RT,RA,-IM		9-12
sync		Synchronization. All instructions that precede sync complete before any instructions that follow sync begin. When sync completes, all storage accesses initiated prior to sync will have completed.		9-182
tlbia		All TLB entries are invalidated and become unavailable for translation by clearing the valid (V) bit in the TLBHI portion of each TLB entry. The rest of the TLB fields unmodified.		9-183
tlbre	RT, RA, WS	If WS = 0: Load TLBHI of the selected TLB entry into RT. Load PID with the contents of the TID field of the selected TLB entry. (RT) $\leftarrow TLBHI[(RA)]$ (PID) $\leftarrow TLB[(RA)]_{TID}$ If WS = 1: Load TLBLO portion of the selected TLB entry into RT. (RT) $\leftarrow TLBLO[(RA)]$		9-184

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tlbrehi	RT, RA	Load TLBHI of the selected TLB entry into RT. Load PID with the contents of the TID field of the selected TLB entry. (RT) \leftarrow TLBHI[(RA)] (PID) \leftarrow TLB[(RA)] _{TID} <i>Extended mnemonic for</i> tlbre RT,RA,0		9-184
tlbrelo	RT, RA	Load TLBLO of the selected TLB entry into RT. (RT) \leftarrow TLBLO[(RA)] <i>Extended mnemonic for</i> tlbre RT,RA,1		9-184
tlbsx	RT, RA, RB	Search the TLB for a valid entry that translates the EA. EA = (RA)0 + (RB). If found, (RT) \leftarrow Index of TLB entry. If not found, (RT) Undefined.		9-186
tlbsx.		If found, (RT) \leftarrow Index of TLB entry. CR[CR0] _{EQ} \leftarrow 1. If not found, (RT) Undefined. CR[CR0] _{EQ} \leftarrow 1.	CR[CR0] _{LT,GT,SO}	
tlbsync		tlbsync does not complete until all previous TLB-update instructions executed by this processor have been received and completed by all other processors. For the PPC405 core, tlbsync is a no-op.		9-187
tlbwe	RS, RA, WS	If WS = 0: Write TLBHI of the selected TLB entry from RS. Write the TID field of the selected TLB entry from the PID register. TLBHI[(RA)] \leftarrow (RS) TLB[(RA)] _{TID} \leftarrow (PID) _{24:31} If WS = 1: Write TLBLO portion of the selected TLB entry from RS. TLBLO[(RA)] \leftarrow (RS)		9-188
tlbwehi	RS, RA	Write TLBHI of the selected TLB entry from RS. Write the TID field of the selected TLB entry from the PID register. TLBHI[(RA)] \leftarrow (RS) TLB[(RA)] _{TID} \leftarrow (PID) _{24:31} <i>Extended mnemonic for</i> tlbwe RS,RA,0		9-188
tlbwelo	RS, RA	Write TLBLO of the selected TLB entry from RS. TLBLO[(RA)] \leftarrow (RS) <i>Extended mnemonic for</i> tlbwe RS,RA,1		9-188

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
trap		Trap unconditionally. <i>Extended mnemonic for</i> tw 31,0,0		9-190
tweq	RA, RB	Trap if (RA) equal to (RB). <i>Extended mnemonic for</i> tw 4,RA,RB		
twge		Trap if (RA) greater than or equal to (RB). <i>Extended mnemonic for</i> tw 12,RA,RB		
twgt		Trap if (RA) greater than (RB). <i>Extended mnemonic for</i> tw 8,RA,RB		
twle		Trap if (RA) less than or equal to (RB). <i>Extended mnemonic for</i> tw 20,RA,RB		
twlge		Trap if (RA) logically greater than or equal to (RB). <i>Extended mnemonic for</i> tw 5,RA,RB		
twlgt		Trap if (RA) logically greater than (RB). <i>Extended mnemonic for</i> tw 1,RA,RB		
twlle		Trap if (RA) logically less than or equal to (RB). <i>Extended mnemonic for</i> tw 6,RA,RB		
twllt		Trap if (RA) logically less than (RB). <i>Extended mnemonic for</i> tw 2,RA,RB		
twlng		Trap if (RA) logically not greater than (RB). <i>Extended mnemonic for</i> tw 6,RA,RB		
twlnl		Trap if (RA) logically not less than (RB). <i>Extended mnemonic for</i> tw 5,RA,RB		
twlt		Trap if (RA) less than (RB). <i>Extended mnemonic for</i> tw 16,RA,RB		
twne		Trap if (RA) not equal to (RB). <i>Extended mnemonic for</i> tw 24,RA,RB		
twng		Trap if (RA) not greater than (RB). <i>Extended mnemonic for</i> tw 20,RA,RB		
twnl		Trap if (RA) not less than (RB). <i>Extended mnemonic for</i> tw 12,RA,RB		
tw	TO, RA, RB	Trap exception is generated if, comparing (RA) with (RB), any condition specified by TO is true.		9-190

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tweqi	RA, IM	Trap if (RA) equal to EXTS(IM). <i>Extended mnemonic for</i> twi 4,RA,IM		9-193
twgei		Trap if (RA) greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM		
twgti		Trap if (RA) greater than EXTS(IM). <i>Extended mnemonic for</i> twi 8,RA,IM		
twlei		Trap if (RA) less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM		
twlgei		Trap if (RA) logically greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM		
twlgti		Trap if (RA) logically greater than EXTS(IM). <i>Extended mnemonic for</i> twi 1,RA,IM		
twllei		Trap if (RA) logically less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM		
twllti		Trap if (RA) logically less than EXTS(IM). <i>Extended mnemonic for</i> twi 2,RA,IM		
twlngi		Trap if (RA) logically not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM		
twlnli		Trap if (RA) logically not less than EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM		
twlti		Trap if (RA) less than EXTS(IM). <i>Extended mnemonic for</i> twi 16,RA,IM		
twnei		Trap if (RA) not equal to EXTS(IM). <i>Extended mnemonic for</i> twi 24,RA,IM		
twngi		Trap if (RA) not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM		
twnli		Trap if (RA) not less than EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM		
twi	TO, RA, IM	Trap exception is generated if, comparing (RA) with EXTS(IM), any condition specified by TO is true.		9-193
wrttee	RS	Write value of RS ₁₆ to MSR[EE].		9-196
wrtteei	E	Write value of E to MSR[EE].		9-197

Table A-1. PPC405 Instruction Syntax Summary (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
xor	RA, RS, RB	XOR (RS) with (RB). Place result in RA.		9-198
xor.			CR[CR0]	
xori	RA, RS, IM	XOR (RS) with (¹⁶ 0 IM). Place result in RA.		9-199
xoris	RA, RS, IM	XOR (RS) with (IM ¹⁶ 0). Place result in RA.		9-200

A.2 Instructions Sorted by Opcode

All 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 A-44) in bits 0:5. Some instructions also have a secondary opcode field (shown as field XO in Figure A-1 through Figure A-9). PPC405 instructions, sorted by primary and secondary opcode, are listed in Table A-2.

The “Form” indicated in the table refers to the arrangement of valid field combinations within the four-byte instruction. See “Instruction Formats,” on page A-41, 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-2. PPC405 Instructions by Opcode

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
3		D	twi	TO, RA, IM	9-193
4	8	X	mulhhwu mulhhwu.	RT, RA, RB	9-124
4	12 (524)	XO	machhwu machhwu. machhwuo machhwuo.	RT, RA, RB	9-98
4	40	X	mulhww mulhww.	RT, RA, RB	9-123
4	44 (556)	XO	machhw machhw. machhwo machhwo.	RT, RA, RB	9-99
4	46 (558)	XO	nmachhw nmachhw. nmachhwo nmachhwo.	RT, RA, RB	9-135

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
4	76 (588)	XO	machwsu	RT, RA, RB	9-101
			machhwsu.		
			machhwsuo		
			machhwsuo.		
4	108 (620)	XO	machws	RT, RA, RB	9-100
			machhws.		
			machhwsuo		
			machhwsuo.		
4	110 (622)	XO	nmachws	RT, RA, RB	9-136
			nmachhws.		
			nmachhwsuo		
			nmachhwsuo.		
4	136	X	mulchwu	RT, RA, RB	9-122
			mulchwu.		
4	140 (652)	XO	macchwu	RT, RA, RB	9-98
			macchwu.		
			macchwuo		
			macchwuo.		
4	168	X	mulchw	RT, RA, RB	9-121
			mulchw.		
4	172 (684)	XO	macchw	RT, RA, RB	9-95
			macchw.		
			macchwuo		
			macchwuo.		
4	174 (686)	XO	nmacchw	RT, RA, RB	9-133
			nmacchw.		
			nmacchwuo		
			nmacchwuo.		
4	204 (716)	XO	macchwsu	RT, RA, RB	9-97
			macchwsu.		
			macchwsuo		
			macchwsuo.		
4	236 (748)	XO	macchws	RT, RA, RB	9-96
			macchws.		
			macchwso		
			macchwso.		
4	238 (750)	XO	nmacchws	RT, RA, RB	9-134
			nmacchws.		
			nmacchwso		
			nmacchwso.		

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
4	392	X	mullhwu	RT, RA, RB	9-128
			mullhwu.		
4	396 (908)	XO	macldwu	RT, RA, RB	9-106
			macldwu.		
			macldwuo		
			macldwuo.		
4	424	X	mullhw	RT, RA, RB	9-127
			mullhw.		
4	428 (940)	XO	macldw	RT, RA, RB	9-103
			macldw.		
			macldwo		
			macldwo.		
4	430 (942)	XO	nmacldw	RT, RA, RB	9-137
			nmacldw.		
			nmacldwo		
			nmacldwo.		
4	492 (972)	XO	macldws	RT, RA, RB	9-104
			macldws.		
			macldwso		
			macldwso.		
4	460 (1004)	XO	macldwsu	RT, RA, RB	9-105
			macldwsu.		
			macldwsuo		
			macldwsuo.		
4	494 (1006)	XO	nmacldws	RT, RA, RB	9-138
			nmacldws.		
			nmacldwso		
			nmacldwso.		
7		D	mulli	RT, RA, IM	9-129
8		D	subfic	RT, RA, IM	9-179
10		D	cmpli	BF, 0, RA, IM	9-37
11		D	cmpi	BF, 0, RA, IM	9-35
12		D	addic	RT, RA, IM	9-10
13		D	addic.	RT, RA, IM	9-11
14		D	addi	RT, RA, IM	9-9
15		D	addis	RT, RA, IM	9-12
16		B	bc	BO, BI, target	9-20
			bca		
			bcl		
			bcla		
17		SC	sc		9-151

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
18		I	b	target	9-19
			ba		
			bl		
			bla		
19	0	XL	mcrf	BF, BFA	9-107
19	16	XL	bclr	BO, BI	9-30
			bclrl		
19	33	XL	crnor	BT, BA, BB	9-43
19	50	XL	rfi		9-145
19	51	XL	rfci		9-144
19	129	XL	crandc	BT, BA, BB	9-40
19	150	XL	isync		9-70
19	193	XL	crxor	BT, BA, BB	9-46
19	225	XL	crnand	BT, BA, BB	9-42
19	257	XL	crand	BT, BA, BB	9-39
19	289	XL	creqv	BT, BA, BB	9-41
19	417	XL	crorc	BT, BA, BB	9-45
19	449	XL	cror	BT, BA, BB	9-44
19	528	XL	bcctr	BO, BI	9-26
			bcctrl		
20		M	rlwimi	RA, RS, SH, MB, ME	9-146
			rlwimi.		
21		M	rlwinm	RA, RS, SH, MB, ME	9-147
			rlwinm.		
23		M	rlwnm	RA, RS, RB, MB, ME	9-150
			rlwnm.		
24		D	ori	RA, RS, IM	9-142
25		D	oris	RA, RS, IM	9-143
26		D	xori	RA, RS, IM	9-199
27		D	xoris	RA, RS, IM	9-200
28		D	andi.	RA, RS, IM	9-17
29		D	andis.	RA, RS, IM	9-18
31	0	X	cmp	BF, 0, RA, RB	9-34
31	4	X	tw	TO, RA, RB	9-190
31	8 (520)	XO	subfc	RT, RA, RB	9-177
			subfc.		
			subfco		
			subfco.		

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	10 (522)	XO	addc	RT, RA, RB	9-7
			addc.		
			addco		
			addco.		
31	11	XO	mulhwu	RT, RA, RB	9-126
			mulhwu.		
31	19	X	mfcr	RT	9-109
31	20	X	lwarx	RT, RA, RB	9-89
31	23	X	lwzcx	RT, RA, RB	9-94
31	24	X	slw	RA, RS, RB	9-169
			slw.		
31	26	X	cntlzw	RA, RS	9-38
			cntlzw.		
31	28	X	and	RA, RS, RB	9-15
			and.		
31	32	X	cmpl	BF, 0, RA, RB	9-36
31	40 (552)	XO	subf	RT, RA, RB	9-176
			subf.		
			subfo		
			subfo.		
31	54	X	dcbst	RA, RB	9-51
31	55	X	lwzux	RT, RA, RB	9-93
31	60	X	andc	RA, RS, RB	9-16
			andc.		
31	75	XO	mulhw	RT, RA, RB	9-125
			mulhw.		
31	83	X	mfmsr	RT	9-111
31	86	X	dcbf	RA, RB	9-49
31	87	X	lbzcx	RT, RA, RB	9-74
31	104 (616)	XO	neg	RT, RA	9-132
			neg.		
			nego		
			nego.		
31	119	X	lbzux	RT, RA, RB	9-73
31	124	X	nor	RA, RS, RB	9-139
			nor.		
31	131	X	wrttee	RS	9-196
31	136 (648)	XO	subfe	RT, RA, RB	9-178
			subfe.		
			subfeo		
			subfeo.		

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	138 (650)	XO	adde	RT, RA, RB	9-8
			adde.		
			addeo		
			addeo.		
31	144	XFX	mtrcf	FXM, RS	9-116
31	146	X	mtmsr	RS	9-118
31	150	X	stwcx.	RS, RA, RB	9-171
31	151	X	stwx	RS, RA, RB	9-175
31	163	X	wrtnei	E	9-197
31	183	X	stwux	RS, RA, RB	9-174
31	200 (712)	XO	subfze	RT, RA, RB	9-181
			subfze.		
			subfzeo		
			subfzeo.		
31	202 (714)	XO	addze	RT, RA	9-14
			addze.		
			addzeo		
			addzeo.		
31	215	X	stbx	RS, RA, RB	9-159
31	232 (744)	XO	subfme	RT, RA, RB	9-180
			subfme.		
			subfmeo		
			subfmeo.		
31	234 (746)	XO	addme	RT, RA	9-13
			addme.		
			addmeo		
			addmeo.		
31	235 (747)	XO	mullw	RT, RA, RB	9-130
			mullw.		
			mullwo		
			mullwo.		
31	246	X	dcbtst	RA, RB	9-53
31	247	X	stbux	RS, RA, RB	9-158
31	262	X	icbt	RA, RB	9-66
31	266 (778)	XO	add	RT, RA, RB	9-6
			add.		
			addo		
			addo.		
31	278	X	dcbt	RA, RB	9-52
31	279	X	lhzx	RT, RA, RB	9-83

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	284	X	eqv	RA, RS, RB	9-62
			eqv.		
31	311	X	lhzux	RT, RA, RB	9-82
31	316	X	xor	RA, RS, RB	9-198
			xor.		
31	323	AFX	mfdcr	RT, DCRN	9-110
31	339	AFX	mfspr	RT, SPRN	9-112
31	343	X	lhax	RT, RA, RB	9-78
31	370	X	tlbia		9-183
31	371	AFX	mftb	RT, TBRN	9-114
31	375	X	lhaux	RT, RA, RB	9-77
31	407	X	sthx	RS, RA, RB	9-164
31	412	X	orc	RA, RS, RB	9-141
			orc.		
31	439	X	sthux	RS, RA, RB	9-163
31	444	X	or	RA, RS, RB	9-140
			or.		
31	451	AFX	mtdcr	DCRN, RS	9-117
31	454	X	dccci	RA, RB	9-56
31	459 (971)	XO	divwu	RT, RA, RB	9-60
			divwu.		
			divwuo		
			divwuo.		
31	467	AFX	mtspr	SPRN, RS	9-119
31	470	X	dcbi	RA, RB	9-50
31	476	X	nand	RA, RS, RB	9-131
			nand.		
31	486	X	dcread	RT, RA, RB	9-57
31	491 (1003)	XO	divw	RT, RA, RB	9-59
			divw.		
			divwo		
			divwo.		
31	512	X	mcrxr	BF	9-108
31	533	X	lswx	RT, RA, RB	9-87
31	534	X	lwbrx	RT, RA, RB	9-79
31	536	X	srw	RA, RS, RB	9-155
			srw.		
31	566	X	tlbsync		9-187
31	597	X	lswi	RT, RA, NB	9-82
31	598	X	sync		9-182
31	661	X	stswx	RS, RA, RB	9-167

Table A-2. PPC405 Instructions by Opcode (continued)

Primary Opcode	Secondary Opcode	Form	Mnemonic	Operands	Page
31	662	X	stwbrx	RS, RA, RB	9-170
31	725	X	stswi	RS, RA, NB	9-166
31	758	X	dcba	RA, RB	9-47
31	790	X	lhbrx	RT, RA, RB	9-79
31	792	X	sraw	RA, RS, RB	9-153
			sraw.		
31	824	X	srawi	RA, RS, SH	9-154
			srawi.		
31	854	X	eiemo		9-61
31	914	X	tlbsx	RT, RA, RB	9-186
			tlbsx.		
31	918	X	sthbrx	RS, RA, RB	9-161
31	922	X	extsh	RA, RS	9-64
			extsh.		
31	946	X	tlbre	RT, RA, WS	9-184
31	954	X	extsb	RA, RS	9-63
			extsb.		
31	966	X	iccci	RA, RB	9-67
31	978	X	tlbwe	RS, RA, WS	9-188
31	982	X	icbi	RA, RB	9-65
31	998	X	icread	RA, RB	9-68
31	1014	X	dcbz	RA, RB	9-54
32		D	lwz	RT, D(RA)	9-91
33		D	lwzu	RT, D(RA)	9-92
34		D	lbz	RT, D(RA)	9-71
35		D	lbzu	RT, D(RA)	9-72
36		D	stw	RS, D(RA)	9-169
37		D	stwu	RS, D(RA)	9-173
38		D	stb	RS, D(RA)	9-156
39		D	stbu	RS, D(RA)	9-157
40		D	lhz	RT, D(RA)	9-80
41		D	lhzu	RT, D(RA)	9-81
42		D	lha	RT, D(RA)	9-75
43		D	lhau	RT, D(RA)	9-76
44		D	sth	RS, D(RA)	9-160
45		D	sthu	RS, D(RA)	9-162
46		D	lmw	RT, D(RA)	9-84
47		D	stmw	RS, D(RA)	9-165

A.3 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 PPC405 core executes all invalid instruction forms without causing an illegal instruction exception.

A.3.1 Instruction Fields

PPC405 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 mcrf instruction.
BFA (11:13)	Specifies a field in the CR used as a source in a mcrf 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 “BO Field on Conditional Branches” on page 2-25.
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 twos-complement integer displacement for load/store instructions.
DCRN (11:20)	Specifies a device control register (DCR).
FXM (12:19)	Field mask used to identify CR fields to be updated by the mtrf 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.
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 “Condition Register (CR)” on page 2-10 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).
TO (6:10)	Specifies the conditions on which to trap, as described under tw and twi instructions.
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.3.2 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-2 on page A-33 indicates which “form” is utilized by each PPC405 opcode. Fields indicated by slashes (/, //, or ///) are reserved. The figures are adapted from the PowerPC User Instruction Set Architecture.

A.3.2.1 I-Form




Figure A-1. I Instruction Format

A.3.2.2 B-Form




Figure A-2. B Instruction Format

A.3.2.3 SC-Form



Figure A-3. SC Instruction Format

A.3.2.4 D-Form




Figure A-4. D Instruction Format

A.3.2.5 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	//		///		Rc	
OPCD	BF	//		///		///		XO		/
OPCD	BF	//		///		U		XO		Rc
OPCD	BF	//		///		///		XO		/
OPCD	TO		RA		RB		XO		/	
OPCD	BT		///		///		XO		Rc	
OPCD	///		RA		RB		XO		/	
OPCD	///		///		///		XO		/	
OPCD	///		///		E	//		XO		/
0	6	11	16	21	26	31				

Figure A-5. X Instruction Format

A.3.2.6 XL-Form

OPCD	BT		BA		BB		XO		/
OPCD	BC		BI		///		XO		LK
OPCD	BF	//		BFA	//		///		/
OPCD	///		///		///		XO		/
0	6	11	16	21	26	31			

Figure A-6. XL Instruction Format

A.3.2.7 XFX-Form

OPCD	RT	SPRF		XO	/
OPCD	RT	DCRF		XO	/
OPCD	RT	/	FXM	/	/
OPCD	RS	SPRF		XO	/
OPCD	RS	DCRF		XO	/
0	6	11	16	21	31

Figure A-7. XFX Instruction Format

A.3.2.8 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.3.2.9 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

Appendix B. Instructions by Category

Chapter 9, “Instruction Set,” contains detailed descriptions of the instructions, their operands, and notation.

Table B-1 summarizes the instruction categories in the PPC405 instruction set. The instructions within each category are listed in subsequent tables.

Table B-1. PPC405 Instruction Set Categories

Storage Reference	load, store
Arithmetic and Logical	add, subtract, negate, multiply, divide, and, andc, or, orc, xor, nand, nor, xnor, sign extension, count leading zeros, multiply accumulate
Comparison	compare, compare logical, compare immediate
Branch	branch, branch conditional, branch to LR, branch to CTR
CR Logical	crand, crandc, cror, crorc, crnand, crnor, crxor, crxnor, move CR field
Rotate/Shift	rotate and insert, rotate and mask, shift left, shift right
Cache Control	invalidate, touch, zero, flush, store, read
Interrupt Control	write to external interrupt enable bit, move to/from MSR, return from interrupt, return from critical interrupt
Processor Management	system call, synchronize, trap, move to/from DCRs, move to/from SPRs, move to/from CR

B.1 Implementation-Specific Instructions

To meet the functional requirements of processors for embedded systems and real-time applications, the PPC405 core defines the implementation-specific instructions summarized in Table B-2.

Table B-2. Implementation-specific Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
dccci	RA, RB	Invalidate the data cache congruence class associated with the effective address (EA) (RA 0) + (RB).		9-56
dcread	RT, RA, RB	Read either tag or data information from the data cache congruence class associated with the EA (RA 0) + (RB). Place the results in RT.		9-57
iccci	RA, RB	Invalidate instruction cache.		9-67
icread	RA, RB	Read either tag or data information from the instruction cache congruence class associated with the EA (RA 0) + (RB). Place the results in ICDBDR.		9-67

Table B-2. Implementation-specific Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
macchw	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ signed $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-95
macchw.			CR[CR0]	
macchwo			XER[SO, OV]	
macchwo.			CR[CR0] XER[SO, OV]	
macchws	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ 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 $(RT) \leftarrow temp_{1:32}$		9-96
macchws.			CR[CR0]	
macchwso			XER[SO, OV]	
macchwso.			CR[CR0] XER[SO, OV]	
macchwsu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee^{32}temp_0)$		9-97
macchwsu.			CR[CR0]	
macchwsuo			XER[SO, OV]	
macchwsuo.			CR[CR0] XER[SO, OV]	
macchwu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-98
macchwu.			CR[CR0]	
macchwuo			XER[SO, OV]	
macchwuo.			CR[CR0] XER[SO, OV]	
machhw	RT, RA, RB	$prod_{0:15} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ signed $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-99
machhw.			CR[CR0]	
machhwo			XER[SO, OV]	
machhwo.			CR[CR0] XER[SO, OV]	
machhws	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ 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 $(RT) \leftarrow temp_{1:32}$		9-100
machhws.			CR[CR0]	
machhwsso			XER[SO, OV]	
machhwsso.			CR[CR0] XER[SO, OV]	
machhwsu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee^{32}temp_0)$		9-101
machhwsu.			CR[CR0]	
machhwsuo			XER[SO, OV]	
machhwsuo.			CR[CR0] XER[SO, OV]	

Table B-2. Implementation-specific Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
machhwu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-102
machhwu.			CR[CR0]	
machhwuo			XER[SO, OV]	
machhwuo.			CR[CR0] XER[SO, OV]	
maclhw	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ signed $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-103
maclhw.			CR[CR0]	
maclhwo			XER[SO, OV]	
maclhwo.			CR[CR0] XER[SO, OV]	
maclhws	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ 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 $(RT) \leftarrow temp_{1:32}$		9-104
maclhws.			CR[CR0]	
maclhwso			XER[SO, OV]	
maclhwso.			CR[CR0] XER[SO, OV]	
maclhwsu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow (temp_{1:32} \vee^{32} temp_0)$		9-105
maclhwsu.			CR[CR0]	
maclhwsuo			XER[SO, OV]	
maclhwsuo.			CR[CR0] XER[SO, OV]	
maclhwu	RT, RA, RB	$prod_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ unsigned $temp_{0:32} \leftarrow prod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-106
maclhwu.			CR[CR0]	
maclhwuo			XER[SO, OV]	
maclhwuo.			CR[CR0] XER[SO, OV]	
mulchw	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ signed		9-121
mulchw.			CR[CR0]	
mulchwu	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{0:15}$ unsigned		9-122
mulchwu.			CR[CR0]	
mulhhw	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ signed		9-123
mulhhw.			CR[CR0]	
mulhhwu	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{0:15} \times (RB)_{0:15}$ unsigned		9-124
mulhhwu.			CR[CR0]	
mullhw	RT, RA, RB	$(RT)_{0:31} \leftarrow (RA)_{16:31} \times (RB)_{16:31}$ signed		9-127
mullhw.			CR[CR0]	

Table B-2. Implementation-specific Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mullhwu	RT, RA, RB	$(RT)_{16:31} \leftarrow (RA)_{0:15} \times (RB)_{16:31}$ unsigned		9-128
mullhwu.			CR[CR0]	
nmacchw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15})$ signed $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-133
nmacchw.			CR[CR0]	
nmacchw0			XER[SO, OV]	
nmacchw0.			CR[CR0] XER[SO, OV]	
nmacchws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{0:15})$ 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 $(RT) \leftarrow temp_{1:32}$		9-134
nmacchws.			CR[CR0]	
nmacchwso			XER[SO, OV]	
nmacchwso.			CR[CR0] XER[SO, OV]	
nmachhw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ signed $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-135
nmachhw.			CR[CR0]	
nmachhw0			XER[SO, OV]	
nmachhw0.			CR[CR0] XER[SO, OV]	
nmachhws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{0:15} \times (RB)_{0:15})$ 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 $(RT) \leftarrow temp_{1:32}$		9-136
nmachhws.			CR[CR0]	
nmachhwso			XER[SO, OV]	
nmachhwso.			CR[CR0] XER[SO, OV]	
nmaclhw	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{16:31})$ signed $temp_{0:32} \leftarrow nprod_{0:31} + (RT)$ $(RT) \leftarrow temp_{1:32}$		9-137
nmaclhw.			CR[CR0]	
nmaclhw0			XER[SO, OV]	
nmaclhw0.			CR[CR0] XER[SO, OV]	
nmaclhws	RT, RA, RB	$nprod_{0:31} \leftarrow -((RA)_{16:31} \times (RB)_{16:31})$ 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 $(RT) \leftarrow temp_{1:32}$		9-138
nmaclhws.			CR[CR0]	
nmaclhwso			XER[SO, OV]	
nmaclhwso.			CR[CR0] XER[SO, OV]	

B.2 Instructions in the IBM PowerPC Embedded Environment

To meet the functional requirements of processors for embedded systems and real-time applications, the IBM PowerPC Embedded Environment defines instructions that are not part of the PowerPC Architecture.

Table B-3 summarizes the PPC405 core instructions in the PowerPC Embedded Environment.

Table B-3. Instructions in the IBM PowerPC Embedded Environment

Mnemonic	Operands	Function	Other Registers Changed	Page
dcba	RA, RB	Speculatively establish the data cache block which contains the EA (RA 0) + (RB).		9-47
dcbf	RA, RB	Flush (store, then invalidate) the data cache block which contains the EA (RA 0) + (RB).		9-49
dcbi	RA, RB	Invalidate the data cache block which contains the EA (RA 0) + (RB).		9-50
dcbst	RA, RB	Store the data cache block which contains the EA (RA 0) + (RB).		9-50
dcbt	RA, RB	Load the data cache block which contains the EA (RA 0) + (RB).		9-52
dcbstst	RA, RB	Load the data cache block which contains the EA (RA 0) + (RB).		9-53
dcbz	RA, RB	Zero the data cache block which contains the EA (RA 0) + (RB).		9-54
eieio		Storage synchronization. All loads and stores that precede the eieio instruction complete before any loads and stores that follow the instruction access main storage. Implemented as sync , which is more restrictive.		9-61
icbi	RA, RB	Invalidate the instruction cache block which contains the EA (RA 0) + (RB).		9-65
icbt	RA, RB	Load the instruction cache block which contains the EA (RA 0) + (RB).		9-66
isync		Synchronize execution context by flushing the prefetch queue.		9-70
mfocr	RT, DCRN	Move from DCR to RT, (RT) ← (DCR(DCRN)).		9-110
mfmsr	RT	Move from MSR to RT, (RT) ← (MSR).		9-111
mfmspr	RT, SPRN	Move from SPR to RT, (RT) ← (SPR(SPRN)). Privileged for all SPRs except LR, CTR, TBHU, TBLU, and XER.		9-112

Table B-3. Instructions in the IBM PowerPC Embedded Environment (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mftb	RT	Move the contents of a Time Base Register (TBR) into RT, $TBRN \leftarrow TBRF_{5:9} \parallel TBRF_{0:4}$ $(RT) \leftarrow (TBR(TBRN))$		9-114
mtdcr	DCRN, RS	Move to DCR from RS, $(DCR(DCRN)) \leftarrow (RS)$.		9-117
mtmsr	RS	Move to MSR from RS, $(MSR) \leftarrow (RS)$.		9-118
mtspr	SPRN, RS	Move to SPR from RS, $(SPR(SPRN)) \leftarrow (RS)$. Privileged for all SPRs except LR, CTR, and XER.		9-119
rfci		Return from critical interrupt $(PC) \leftarrow (SRR2)$. $(MSR) \leftarrow (SRR3)$.		9-144
rfi		Return from interrupt. $(PC) \leftarrow (SRR0)$. $(MSR) \leftarrow (SRR1)$.		9-145
tlbia		All of the entries in the TLB are invalidated and become unavailable for translation by clearing the valid (V) bit in the TLBHI portion of each TLB entry. The rest of the fields in the TLB entries are unmodified.		9-183
tlbre	RT, RA, WS	If WS = 0: Load TLBHI portion of the selected TLB entry into RT. Load the PID register with the contents of the TID field of the selected TLB entry. $(RT) \leftarrow TLBHI[(RA)]$ $(PID) \leftarrow TLB[(RA)]_{TID}$ If WS = 1: Load TLBLO portion of the selected TLB entry into RT. $(RT) \leftarrow TLBLO[(RA)]$		9-184

Table B-3. Instructions in the IBM PowerPC Embedded Environment (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tlbsx	RT,RA,RB	Search the TLB array for a valid entry which translates the EA EA = (RA 0) + (RB). If found, (RT) ← Index of TLB entry. If not found, (RT) Undefined.		9-186
tlbsx.		If found, (RT) ← Index of TLB entry. CR[CR0] _{EQ} ← 1. If not found, (RT) Undefined. CR[CR0] _{EQ} ← 1.	CR[CR0] _{LT,GT,SO}	
tlbsync		tlbsync does not complete until all previous TLB-update instructions executed by this processor have been received and completed by all other processors. For the PPC405 core, tlbsync is a no-op.		9-187
tlbwe	RS, RA,WS	If WS = 0: Write TLBHI portion of the selected TLB entry from RS. Write the TID field of the selected TLB entry from the PID register. TLBHI[(RA)] ← (RS) TLB[(RA)] _{TID} ← (PID) _{24:31} If WS = 1: Write TLBLO portion of the selected TLB entry from RS. TLBLO[(RA)] ← (RS)		9-188
wrtee	RS	Write value of RS ₁₆ to MSR[EE].		9-196
wrteei	E	Write value of E to MSR[EE].		9-197

B.3 Privileged Instructions

Table B-4 lists instructions that are under control of the MSR[PR] bit. These instructions are not allowed to be executed when MSR[PR] = 1:

Table B-4. Privileged Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
dcbi	RA, RB	Invalidate the data cache block which contains the EA (RA 0) + (RB).		9-50
dccci	RA, RB	Invalidate the data cache congruence class associated with the EA (RA 0) + (RB).		9-56

Table B-4. Privileged Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
dcread	RT, RA, RB	Read either tag or data information from the data cache congruence class associated with the EA (RA 0) + (RB). Place the results in RT.		9-57
iccci	RA, RB	Invalidate instruction cache.		9-67
icread	RA, RB	Read either tag or data information from the instruction cache congruence class associated with the EA (RA 0) + (RB). Place the results in ICDBDR.		9-68
mfdcr	RT, DCRN	Move from DCR to RT, (RT) ← (DCR(DCRN)).		9-110
mfmsr	RT	Move from MSR to RT, (RT) ← (MSR).		9-111
mfspr	RT, SPRN	Move from SPR to RT, (RT) ← (SPR(SPRN)). Privileged for all SPRs except LR, CTR, TBHU, TBLU, and XER.		9-112
mtdcr	DCRN, RS	Move to DCR from RS, (DCR(DCRN)) ← (RS).		9-117
mtmsr	RS	Move to MSR from RS, (MSR) ← (RS).		9-118
mtspr	SPRN, RS	Move to SPR from RS, (SPR(SPRN)) ← (RS). Privileged for all SPRs except LR, CTR, and XER.		9-119
rfdi		Return from critical interrupt (PC) ← (SRR2). (MSR) ← (SRR3).		9-144
rfdi		Return from interrupt. (PC) ← (SRR0). (MSR) ← (SRR1).		9-145
tlbre	RT, RA, WS	If WS = 0: Load TLBHI portion of the selected TLB entry into RT. Load the PID register with the contents of the TID field of the selected TLB entry. (RT) ← TLBHI[(RA)] (PID) ← TLB[(RA)] _{TID} If WS = 1: Load TLBLO portion of the selected TLB entry into RT. (RT) ← TLBLO[(RA)]		9-184

Table B-4. Privileged Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tlbsx	RT,RA,RB	Search the TLB array for a valid entry which translates the EA EA = (RA 0) + (RB). If found, (RT) ← Index of TLB entry. If not found, (RT) Undefined.		9-186
tlbsx.		If found, (RT) ← Index of TLB entry. CR[CR0] _{EQ} ← 1. If not found, (RT) Undefined. CR[CR0] _{EQ} ← 1.	CR[CR0] _{LT,GT,SO}	
tlbwe	RS, RA,WS	If WS = 0: Write TLBHI portion of the selected TLB entry from RS. Write the TID field of the selected TLB entry from the PID register. TLBHI[(RA)] ← (RS) TLB[(RA)] _{TID} ← (PID) _{24:31} If WS = 1: Write TLBLO portion of the selected TLB entry from RS. TLBLO[(RA)] ← (RS)		9-188
wrtee	RS	Write value of RS ₁₆ to the External Enable bit (MSR[EE]).		9-196
wrteei	E	Write value of E to the External Enable bit (MSR[EE]).		9-197

B.4 Assembler Extended Mnemonics

In the appendix “Assembler Extended Mnemonics” of the PowerPC Architecture, it is required that a PowerPC assembler support at least a minimal set of extended mnemonics. These mnemonics encode to the opcodes of other instructions; the only benefit of extended mnemonics is improved usability. Code using extended mnemonics can be easier to write and to understand. Table B-5 lists the extended mnemonics required for the PPC405.

Note for every Branch Conditional mnemonic:

Bit 4 of the BO field provides a hint about the most likely outcome of a conditional branch. (“Branch Prediction” on page 2-26 describes branch prediction). Assemblers should set BO₄ = 0 unless a specific reason exists otherwise. In the BO field values specified in the following table, BO₄ = 0 has always been assumed. The assembler must allow the programmer to specify branch prediction. To do this, the assembler will support a suffix to every conditional branch mnemonic, as follows:

- + Predict branch to be taken.
- Predict branch not to be taken.

As specific examples, **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 2-26).

Table B-5. Extended Mnemonics for PPC405

Mnemonic	Operands	Function	Other Registers Changed	Page
bctr		Branch unconditionally to address in CTR. <i>Extended mnemonic for</i> bcctr 20,0		9-26
bctrl		<i>Extended mnemonic for</i> bcctrl 20,0	(LR) ← CIA + 4	
bdnz	target	Decrement CTR. Branch if CTR ≠ 0. <i>Extended mnemonic for</i> bc 16,0,target		9-20
bdnza		<i>Extended mnemonic for</i> bca 16,0,target		
bdnzl		<i>Extended mnemonic for</i> bcl 16,0,target	(LR) ← CIA + 4.	
bdnzla		<i>Extended mnemonic for</i> bcla 16,0,target	(LR) ← CIA + 4.	
bdnzlr		Decrement CTR. Branch, if CTR ≠ 0, to address in LR. <i>Extended mnemonic for</i> bclr 16,0		
bdnzlrl		<i>Extended mnemonic for</i> bclrl 16,0	(LR) ← CIA + 4.	9-30
bdnzf	cr_bit, target	Decrement CTR. Branch if CTR ≠ 0 AND $CR_{cr_bit} = 0$. <i>Extended mnemonic for</i> bc 0,cr_bit,target		
bdnzfa		<i>Extended mnemonic for</i> bca 0,cr_bit,target		
bdnzfl		<i>Extended mnemonic for</i> bcl 0,cr_bit,target	(LR) ← CIA + 4.	
bdnzfla		<i>Extended mnemonic for</i> bcla 0,cr_bit,target	(LR) ← CIA + 4.	
bdnzflr	cr_bit	Decrement CTR. Branch, if CTR ≠ 0 AND $CR_{cr_bit} = 0$, to address in LR. <i>Extended mnemonic for</i> bclr 0,cr_bit		9-30
bdnzflrl		<i>Extended mnemonic for</i> bclrl 0,cr_bit	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bdnzt	cr_bit, target	Decrement CTR. Branch if CTR \neq 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 8,cr_bit,target		9-20
bdnzta		<i>Extended mnemonic for</i> bca 8,cr_bit,target		
bdnztl		<i>Extended mnemonic for</i> bcl 8,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnztla		<i>Extended mnemonic for</i> bcla 8,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdnztlr	cr_bit	Decrement CTR. Branch, if CTR \neq 0 AND CR _{cr_bit} = 1, to address in LR. <i>Extended mnemonic for</i> bclr 8,cr_bit		9-30
bdnztlrl		<i>Extended mnemonic for</i> bclrl 8,cr_bit	(LR) \leftarrow CIA + 4.	
bdz	target	Decrement CTR. Branch if CTR = 0. <i>Extended mnemonic for</i> bc 18,0,target		9-20
bdza		<i>Extended mnemonic for</i> bca 18,0,target		
bdzl		<i>Extended mnemonic for</i> bcl 18,0,target	(LR) \leftarrow CIA + 4.	
bdzla		<i>Extended mnemonic for</i> bcla 18,0,target	(LR) \leftarrow CIA + 4.	
bdzlr		Decrement CTR. Branch, if CTR = 0, to address in LR. <i>Extended mnemonic for</i> bclr 18,0		9-30
bdzlrl		<i>Extended mnemonic for</i> bclrl 18,0	(LR) \leftarrow CIA + 4.	
bdzfb	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 2,cr_bit,target		9-20
bdzfa		<i>Extended mnemonic for</i> bca 2,cr_bit,target		
bdzfl		<i>Extended mnemonic for</i> bcl 2,cr_bit,target	(LR) \leftarrow CIA + 4.	
bdzfla		<i>Extended mnemonic for</i> bcla 2,cr_bit,target	(LR) \leftarrow CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bdzflr	cr_bit	Decrement CTR. Branch, if CTR = 0 AND CR _{cr_bit} = 0 to address in LR. <i>Extended mnemonic for</i> bclr 2,cr_bit		9-30
bdzflrl		<i>Extended mnemonic for</i> bclrl 2,cr_bit	(LR) ← CIA + 4.	
bdzt	cr_bit, target	Decrement CTR. Branch if CTR = 0 AND CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 10,cr_bit,target		9-20
bdzta		<i>Extended mnemonic for</i> bca 10,cr_bit,target		
bdztl		<i>Extended mnemonic for</i> bcl 10,cr_bit,target	(LR) ← CIA + 4.	
bdztlr		<i>Extended mnemonic for</i> bcla 10,cr_bit,target	(LR) ← CIA + 4.	
bdztlr	cr_bit	Decrement CTR. Branch, if CTR = 0 AND CR _{cr_bit} = 1, to address in LR. <i>Extended mnemonic for</i> bclr 10,cr_bit		9-30
bdztlrl		<i>Extended mnemonic for</i> bclrl 10,cr_bit	(LR) ← CIA + 4.	
beq	[cr_field,] target	Branch if equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+2,target		9-20
beqa		<i>Extended mnemonic for</i> bca 12,4*cr_field+2,target		
beql		<i>Extended mnemonic for</i> bcl 12,4*cr_field+2,target	(LR) ← CIA + 4.	
beqla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+2,target	(LR) ← CIA + 4.	
beqctr	[cr_field]	Branch, if equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+2		9-26
beqctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+2	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
beqlr	[cr_field]	Branch, if equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+2		9-30
beqlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+2	(LR) ← CIA + 4.	
bf	cr_bit, target	Branch if CR _{cr_bit} = 0. <i>Extended mnemonic for</i> bc 4,cr_bit,target		9-20
bfa		<i>Extended mnemonic for</i> bca 4,cr_bit,target		
bfl		<i>Extended mnemonic for</i> bcl 4,cr_bit,target	(LR) ← CIA + 4.	
bfla		<i>Extended mnemonic for</i> bcla 4,cr_bit,target	(LR) ← CIA + 4.	
bfctr	cr_bit	Branch, if CR _{cr_bit} = 0, to address in CTR. <i>Extended mnemonic for</i> bcctr 4,cr_bit		9-26
bfctrl		<i>Extended mnemonic for</i> bcctrl 4,cr_bit	(LR) ← CIA + 4.	
bflr	cr_bit	Branch, if CR _{cr_bit} = 0, to address in LR. <i>Extended mnemonic for</i> bclr 4,cr_bit		9-30
bflrl		<i>Extended mnemonic for</i> bclrl 4,cr_bit	(LR) ← CIA + 4.	
bge	[cr_field,] target	Branch if greater than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target		9-20
bgea		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target		
bgel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bgela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bgectr	[cr_field]	Branch, if greater than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0		9-26
bgectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bgehr	[cr_field]	Branch, if greater than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0		9-30
bgehrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bgt	[cr_field,] target	Branch if greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+1,target		9-20
bgta		<i>Extended mnemonic for</i> bca 12,4*cr_field+1,target		
bgtl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+1,target	(LR) ← CIA + 4.	
bgvla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+1,target	(LR) ← CIA + 4.	
bgctr	[cr_field]	Branch, if greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+1		9-26
bgctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+1	(LR) ← CIA + 4.	
bgtlr	[cr_field]	Branch, if greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+1		9-30
bgtlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+1	(LR) ← CIA + 4.	
ble	[cr_field,] target	Branch if less than or equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target		9-20
blea		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target		
blel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	(LR) ← CIA + 4.	
blela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
blectr	[cr_field]	Branch, if less than or equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1		9-26
blectrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.	
blelr	[cr_field]	Branch, if less than or equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1		9-30
blelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) ← CIA + 4.	
blr		Branch, unconditionally, to address in LR. <i>Extended mnemonic for</i> bclr 20,0		9-30
blrl		<i>Extended mnemonic for</i> bclrl 20,0	(LR) ← CIA + 4.	
blt	[cr_field,] target	Branch if less than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+0,target		9-20
blta		<i>Extended mnemonic for</i> bca 12,4*cr_field+0,target		
bltl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+0,target	(LR) ← CIA + 4.	
bltla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+0,target	(LR) ← CIA + 4.	
bltctr	[cr_field]	Branch, if less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+0		9-26
bltctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+0	(LR) ← CIA + 4.	
bltlr	[cr_field]	Branch, if less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+0		9-30
bltlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+0	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bne	[cr_field,] target	Branch if not equal. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+2,target		9-20
bnea		<i>Extended mnemonic for</i> bca 4,4*cr_field+2,target		
bnel		<i>Extended mnemonic for</i> bcl 4,4*cr_field+2,target	(LR) ← CIA + 4.	
bnela		<i>Extended mnemonic for</i> bcla 4,4*cr_field+2,target	(LR) ← CIA + 4.	
bnctr	[cr_field]	Branch, if not equal, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+2		9-26
bnctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+2	(LR) ← CIA + 4.	
bnelr	[cr_field]	Branch, if not equal, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+2		9-30
bnelrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+2	(LR) ← CIA + 4.	
bng	[cr_field,] target	Branch, if not greater than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+1,target		9-20
bnga		<i>Extended mnemonic for</i> bca 4,4*cr_field+1,target		
bngl		<i>Extended mnemonic for</i> bcl 4,4*cr_field+1,target	(LR) ← CIA + 4.	
bngla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+1,target	(LR) ← CIA + 4.	
bngctr	[cr_field]	Branch, if not greater than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+1		9-26
bngctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+1	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bnglr	[cr_field]	Branch, if not greater than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+1		9-30
bnglrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+1	(LR) ← CIA + 4.	
bnl	[cr_field,] target	Branch if not less than. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+0,target		9-20
bnla		<i>Extended mnemonic for</i> bca 4,4*cr_field+0,target		
bnll		<i>Extended mnemonic for</i> bcl 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bnlla		<i>Extended mnemonic for</i> bcla 4,4*cr_field+0,target	(LR) ← CIA + 4.	
bnlctr	[cr_field]	Branch, if not less than, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+0		9-26
bnlctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bnllr	[cr_field]	Branch, if not less than, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+0		9-30
bnllrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+0	(LR) ← CIA + 4.	
bns	[cr_field,] target	Branch if not summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target		9-20
bnsa		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target		
bnsi		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnsia		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bnsctr	[cr_field]	Branch, if not summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3		9-26
bnsctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnsldr	[cr_field]	Branch, if not summary overflow, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3		9-30
bnsldr		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnu	[cr_field,] target	Branch if not unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 4,4*cr_field+3,target		9-20
bnuar		<i>Extended mnemonic for</i> bca 4,4*cr_field+3,target		
bnul		<i>Extended mnemonic for</i> bcl 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnular		<i>Extended mnemonic for</i> bcla 4,4*cr_field+3,target	(LR) ← CIA + 4.	
bnuctr	[cr_field]	Branch, if not unordered, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 4,4*cr_field+3		9-26
bnuctrl		<i>Extended mnemonic for</i> bcctrl 4,4*cr_field+3	(LR) ← CIA + 4.	
bnulr	[cr_field]	Branch, if not unordered, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 4,4*cr_field+3		9-30
bnulrl		<i>Extended mnemonic for</i> bclrl 4,4*cr_field+3	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bso	[cr_field,] target	Branch if summary overflow. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target		9-20
bsoa		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target		
bsol		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bsola		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bsoctr	[cr_field]	Branch, if summary overflow, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3		9-26
bsoctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) ← CIA + 4.	
bsolr	[cr_field]	Branch, if summary overflow, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3		9-30
bsolrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) ← CIA + 4.	
bt	cr_bit, target	Branch if CR _{cr_bit} = 1. <i>Extended mnemonic for</i> bc 12,cr_bit,target		9-20
bta		<i>Extended mnemonic for</i> bca 12,cr_bit,target		
btl		<i>Extended mnemonic for</i> bcl 12,cr_bit,target	(LR) ← CIA + 4.	
btla		<i>Extended mnemonic for</i> bcla 12,cr_bit,target	(LR) ← CIA + 4.	
btctr	cr_bit	Branch if CR _{cr_bit} = 1, to address in CTR. <i>Extended mnemonic for</i> bcctr 12,cr_bit		9-26
btctrl		<i>Extended mnemonic for</i> bcctrl 12,cr_bit	(LR) ← CIA + 4.	
btlr	cr_bit	Branch, if CR _{cr_bit} = 1, to address in LR. <i>Extended mnemonic for</i> bclr 12,cr_bit		9-30
btlrl		<i>Extended mnemonic for</i> bclrl 12,cr_bit	(LR) ← CIA + 4.	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
bun	[cr_field,] target	Branch if unordered. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bc 12,4*cr_field+3,target		9-20
buna		<i>Extended mnemonic for</i> bca 12,4*cr_field+3,target		
bunl		<i>Extended mnemonic for</i> bcl 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bunla		<i>Extended mnemonic for</i> bcla 12,4*cr_field+3,target	(LR) ← CIA + 4.	
bunctr	[cr_field]	Branch, if unordered, to address in CTR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bcctr 12,4*cr_field+3		9-26
bunctrl		<i>Extended mnemonic for</i> bcctrl 12,4*cr_field+3	(LR) ← CIA + 4.	
bunlr	[cr_field]	Branch, if unordered, to address in LR. Use CR0 if cr_field is omitted. <i>Extended mnemonic for</i> bclr 12,4*cr_field+3		9-30
bunlrl		<i>Extended mnemonic for</i> bclrl 12,4*cr_field+3	(LR) ← CIA + 4.	
clrlwi	RA, RS, n	Clear left immediate. (n < 32) $(RA)_{0:n-1} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,n,31		9-147
clrlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,n,31	CR[CR0]	
clrlslwi	RA, RS, b, n	Clear left and shift left immediate. (n ≤ b < 32) $(RA)_{b-n:31-n} \leftarrow (RS)_{b:31}$ $(RA)_{32-n:31} \leftarrow {}^n0$ $(RA)_{0:b-n-1} \leftarrow {}^{b-n}0$ <i>Extended mnemonic for</i> rlwinm RA,RS,n,b-n,31-n		9-147
clrlslwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,b-n,31-n	CR[CR0]	
clrrwi	RA, RS, n	Clear right immediate. (n < 32) $(RA)_{32-n:31} \leftarrow {}^n0$ <i>Extended mnemonic for</i> rlwinm RA,RS,0,0,31-n		9-147
clrrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,0,0,31-n	CR[CR0]	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
cmplw	[BF,] RA, RB	Compare Logical Word. Use CR0 if BF is omitted. <i>Extended mnemonic for cmpl BF,0,RA,RB</i>		9-36
cmplwi	[BF,] RA, IM	Compare Logical Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for cmpli BF,0,RA,IM</i>		9-37
cmpw	[BF,] RA, RB	Compare Word. Use CR0 if BF is omitted. <i>Extended mnemonic for cmp BF,0,RA,RB</i>		9-34
cmpwi	[BF,] RA, IM	Compare Word Immediate. Use CR0 if BF is omitted. <i>Extended mnemonic for cmpi BF,0,RA,IM</i>		9-35
crclr	bx	Condition register clear. <i>Extended mnemonic for crxor bx,bx,bx</i>		9-46
crmove	bx, by	Condition register move. <i>Extended mnemonic for cror bx,by,by</i>		9-44
crnot	bx, by	Condition register not. <i>Extended mnemonic for crnor bx,by,by</i>		9-43
crset	bx	Condition register set. <i>Extended mnemonic for creqv bx,bx,bx</i>		9-41
extlwi	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 rlwinm RA,RS,b,0,n-1</i>		9-147
extlwi.		<i>Extended mnemonic for rlwinm. RA,RS,b,0,n-1</i>	CR[CR0]	
extrwi	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 rlwinm RA,RS,b+n,32-n,31</i>		9-147
extrwi.		<i>Extended mnemonic for rlwinm. RA,RS,b+n,32-n,31</i>	CR[CR0]	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
inslwi	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> rlwimi RA,RS,32-b,b,b+n-1		9-146
inslwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b,b,b+n-1	CR[CR0]	
insrwi	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> rlwimi RA,RS,32-b-n,b,b+n-1		9-146
insrwi.		<i>Extended mnemonic for</i> rlwimi. RA,RS,32-b-n,b,b+n-1	CR[CR0]	
la	RT, D(RA)	Load address. ($RA \neq 0$) D is an offset from a base address that is assumed to be (RA). $(RT) \leftarrow (RA) + EXTS(D)$ <i>Extended mnemonic for</i> addi RT,RA,D		9-9
li	RT, IM	Load immediate. $(RT) \leftarrow EXTS(IM)$ <i>Extended mnemonic for</i> addi RT,0,value		9-9
lis	RT, IM	Load immediate shifted. $(RT) \leftarrow (IM \ll 16)$ <i>Extended mnemonic for</i> addis RT,0,value		9-12

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mfccr0 mfctr mfdac1 mfdac2 mfdear mfd bcr0 mfd bcr1 mfd bsr mfd ccr mfd cwr mfd vc1 mfd vc2 mfesr mf evpr mf iac1 mf iac2 mf iac3 mf iac4 mf iccr mf ic bdr mflr mf pid mf pit mf pvr mf sgr mf sler mf sprg0 mf sprg1 mf sprg2 mf sprg3 mf sprg4 mf sprg5 mf sprg6 mf sprg7 mf srr0 mf srr1 mf srr2 mf srr3 mf su0r mf tcr mf tsr mf xer mf zpr	RT	Move from special purpose register (SPR) SPRN. <i>Extended mnemonic for mfspr RT,SPRN</i> See Table 10.5, "Special Purpose Registers," on page 10-2 for listing of valid SPRN values.		9-112
mftb	RT	Move the contents of TBL into RT, (RT) ← (TBL) <i>Extended mnemonic for mftb RT,TBL</i>		9-114

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mftbu	RT	Move the contents of TBU into RT, (RT) ← (TBU) <i>Extended mnemonic for</i> mftb RT,TBU		9-114
mr	RT, RS	Move register. (RT) ← (RS) <i>Extended mnemonic for</i> or RT,RS,RS		9-140
mr.		<i>Extended mnemonic for</i> or. RT,RS,RS	CR[CR0]	
mtr	RS	Move to Condition Register. <i>Extended mnemonic for</i> mtrf 0xFF,RS		9-116

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
mtccr0 mtctr mtdac1 mtdac2 mtdbcr0 mtdbcr1 mtdbsr mtdccr mtdear mtdcwr mtdvc1 mtdvc2 mtesr mtevpr mtiac1 mtiac2 mtiac3 mtiac4 mticcr mticbdr mtlr mtpid mtpit mtpvr mtsgr mtsler mtsprg0 mtsprg1 mtsprg2 mtsprg3 mtsprg4 mtsprg5 mtsprg6 mtsprg7 mtsrr0 mtsrr1 mtsrr2 mtsrr3 mtsu0r mttcr mttsr mtxer mtzpr	RS	Move to SPR SPRN. <i>Extended mnemonic for mtspr SPRN,RS</i> See Table 10.5, "Special Purpose Registers," on page 10-2 for listing of valid SPRN values.		9-119
nop		Preferred no-op; triggers optimizations based on no-ops. <i>Extended mnemonic for ori 0,0,0</i>		9-142

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
not	RA, RS	Complement register. (RA) $\leftarrow \neg$ (RS) <i>Extended mnemonic for</i> nor RA,RS,RS		9-139
not.		<i>Extended mnemonic for</i> nor. RA,RS,RS	CR[CR0]	
rotlw	RA, RS, RB	Rotate left. (RA) \leftarrow ROTL((RS), (RB) _{27:31}) <i>Extended mnemonic for</i> rlwnm RA,RS,RB,0,31		9-150
rotlw.		<i>Extended mnemonic for</i> rlwnm. RA,RS,RB,0,31	CR[CR0]	
rotlwi	RA, RS, n	Rotate left immediate. (RA) \leftarrow ROTL((RS), n) <i>Extended mnemonic for</i> rlwinm RA,RS,n,0,31		9-147
rotlwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31	CR[CR0]	
rotrwi	RA, RS, n	Rotate right immediate. (RA) \leftarrow ROTL((RS), 32-n) <i>Extended mnemonic for</i> rlwinm RA,RS,32-n,0,31		9-147
rotrwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,0,31	CR[CR0]	
slwi	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> rlwinm RA,RS,n,0,31-n		9-147
slwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,n,0,31-n	CR[CR0]	
srwi	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> rlwinm RA,RS,32-n,n,31		9-147
srwi.		<i>Extended mnemonic for</i> rlwinm. RA,RS,32-n,n,31	CR[CR0]	

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
sub	RT, RA, RB	Subtract (RB) from (RA). (RT) \leftarrow \neg (RB) + (RA) + 1. <i>Extended mnemonic for</i> subf RT,RB,RA		9-176
sub.		<i>Extended mnemonic for</i> subf. RT,RB,RA	CR[CR0]	
subo		<i>Extended mnemonic for</i> subfo RT,RB,RA	XER[SO, OV]	
subo.		<i>Extended mnemonic for</i> subfo. RT,RB,RA	CR[CR0] XER[SO, OV]	
subc	RT, RA, RB	Subtract (RB) from (RA). (RT) \leftarrow \neg (RB) + (RA) + 1. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> subfc RT,RB,RA		9-177
subc.		<i>Extended mnemonic for</i> subfc. RT,RB,RA	CR[CR0]	
subco		<i>Extended mnemonic for</i> subfco RT,RB,RA	XER[SO, OV]	
subco.		<i>Extended mnemonic for</i> subfco. RT,RB,RA	CR[CR0] XER[SO, OV]	
subi	RT, RA, IM	Subtract EXTS(IM) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> addi RT,RA,-IM		9-9
subic	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> addic RT,RA,-IM		9-10
subic.	RT, RA, IM	Subtract EXTS(IM) from (RA). Place result in RT. Place carry-out in XER[CA]. <i>Extended mnemonic for</i> addic. RT,RA,-IM	CR[CR0]	9-11
subis	RT, RA, IM	Subtract (IM ¹⁶ 0) from (RA 0). Place result in RT. <i>Extended mnemonic for</i> addis RT,RA,-IM		9-12

Table B-5. Extended Mnemonics for PPC405 (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tweqi	RA, IM	Trap if (RA) equal to EXTS(IM). <i>Extended mnemonic for</i> twi 4,RA,IM		9-190
twgei		Trap if (RA) greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM		
twgti		Trap if (RA) greater than EXTS(IM). <i>Extended mnemonic for</i> twi 8,RA,IM		
twlei		Trap if (RA) less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM		
twlgei		Trap if (RA) logically greater than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM		
twlgti		Trap if (RA) logically greater than EXTS(IM). <i>Extended mnemonic for</i> twi 1,RA,IM		
twllei		Trap if (RA) logically less than or equal to EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM		
twllti		Trap if (RA) logically less than EXTS(IM). <i>Extended mnemonic for</i> twi 2,RA,IM		
twlngi		Trap if (RA) logically not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 6,RA,IM		
twlnli		Trap if (RA) logically not less than EXTS(IM). <i>Extended mnemonic for</i> twi 5,RA,IM		
twlti		Trap if (RA) less than EXTS(IM). <i>Extended mnemonic for</i> twi 16,RA,IM		
twnei		Trap if (RA) not equal to EXTS(IM). <i>Extended mnemonic for</i> twi 24,RA,IM		
twngi		Trap if (RA) not greater than EXTS(IM). <i>Extended mnemonic for</i> twi 20,RA,IM		
twnli		Trap if (RA) not less than EXTS(IM). <i>Extended mnemonic for</i> twi 12,RA,IM		

B.5 Storage Reference Instructions

The PPC405 uses load and store instructions to transfer data between memory and the general purpose registers. Load and store instructions operate on byte, halfword and word data. The storage reference instructions also support loading or storing multiple registers, character strings, and byte-reversed data. Table B-6 shows the storage reference instructions available for use in the PPC405.

Table B-6. Storage Reference Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
lbz	RT, D(RA)	Load byte from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) \leftarrow $^{24}0$ MS(EA,1).		9-71
lbzu	RT, D(RA)	Load byte from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) \leftarrow $^{24}0$ MS(EA,1). Update the base address, (RA) \leftarrow EA.		9-72
lbzux	RT, RA, RB	Load byte from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{24}0$ MS(EA,1). Update the base address, (RA) \leftarrow EA.		9-73
lbzx	RT, RA, RB	Load byte from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{24}0$ MS(EA,1).		9-74
lha	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and sign extend, (RT) \leftarrow EXTS(MS(EA,2)).		9-75
lhau	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and sign extend, (RT) \leftarrow EXTS(MS(EA,2)). Update the base address, (RA) \leftarrow EA.		9-76
lhaux	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and sign extend, (RT) \leftarrow EXTS(MS(EA,2)). Update the base address, (RA) \leftarrow EA.		9-77
lhax	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and sign extend, (RT) \leftarrow EXTS(MS(EA,2)).		9-78
lhbrx	RT, RA, RB	Load halfword from EA = (RA 0) + (RB), then reverse byte order and pad left with zeroes, (RT) \leftarrow $^{16}0$ MS(EA+1,1) MS(EA,1).		9-79
lhz	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) \leftarrow $^{16}0$ MS(EA,2).		9-80

Table B-6. Storage Reference Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
lhzu	RT, D(RA)	Load halfword from EA = (RA 0) + EXTS(D) and pad left with zeroes, (RT) \leftarrow $^{16}0$ MS(EA,2). Update the base address, (RA) \leftarrow EA.		9-80
lhzux	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{16}0$ MS(EA,2). Update the base address, (RA) \leftarrow EA.		9-82
lhzx	RT, RA, RB	Load halfword from EA = (RA 0) + (RB) and pad left with zeroes, (RT) \leftarrow $^{16}0$ MS(EA,2).		9-83
lmw	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).		9-84
lswi	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 $R_{FINAL} \leftarrow ((RT + CEIL(n/4) - 1) \% 32)$. GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R_{FINAL} .		9-85
lswx	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 $R_{FINAL} \leftarrow ((RT + CEIL(n/4) - 1) \% 32)$. GPR(0) is consecutive to GPR(31). RA is not altered unless RA = R_{FINAL} . RB is not altered unless RB = R_{FINAL} . If $n=0$, content of RT is undefined.		9-87
lwarx	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) \leftarrow MS(EA,4). Set the Reservation bit.		9-89
lwbrx	RT, RA, RB	Load word from EA = (RA 0) + (RB) then reverse byte order, (RT) \leftarrow MS(EA+3,1) MS(EA+2,1) MS(EA+1,1) MS(EA,1).		9-90
lwz	RT, D(RA)	Load word from EA = (RA 0) + EXTS(D) and place in RT, (RT) \leftarrow MS(EA,4).		9-91

Table B-6. Storage Reference Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
lwzu	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.		9-92
lwzux	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4). Update the base address, (RA) ← EA.		9-93
lwzx	RT, RA, RB	Load word from EA = (RA 0) + (RB) and place in RT, (RT) ← MS(EA,4).		9-94
stb	RS, D(RA)	Store byte (RS) _{24:31} in memory at EA = (RA 0) + EXTS(D).		9-156
stbu	RS, D(RA)	Store byte (RS) _{24:31} in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) ← EA.		9-157
stbux	RS, RA, RB	Store byte (RS) _{24:31} in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		9-158
stbx	RS, RA, RB	Store byte (RS) _{24:31} in memory at EA = (RA 0) + (RB).		9-159
sth	RS, D(RA)	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + EXTS(D).		9-160
sthbrx	RS, RA, RB	Store halfword (RS) _{16:31} byte-reversed in memory at EA = (RA 0) + (RB). MS(EA, 2) ← (RS) _{24:31} (RS) _{16:23}		9-161
sthu	RS, D(RA)	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) ← EA.		9-162
sthux	RS, RA, RB	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + (RB). Update the base address, (RA) ← EA.		9-163
sthx	RS, RA, RB	Store halfword (RS) _{16:31} in memory at EA = (RA 0) + (RB).		9-164
stmw	RS, D(RA)	Store consecutive words from RS through GPR(31) in memory starting at EA = (RA 0) + EXTS(D).		9-165

Table B-6. Storage Reference Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
stswi	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).		9-166
stswx	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).	9-166	9-167
stw	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXTS(D).	9-166	9-169
stwbrx	RS, RA, RB	Store word (RS) byte-reversed in memory at EA = (RA 0) + (RB). $MS(EA, 4) \leftarrow (RS)_{24:31} \parallel (RS)_{16:23} \parallel (RS)_{8:15} \parallel (RS)_{0:7}$	9-166	9-170
stwcx.	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB) only if the reservation bit is set. if RESERVE = 1 then $MS(EA, 4) \leftarrow (RS)$ RESERVE $\leftarrow 0$ $(CR[CR0]) \leftarrow {}^20 \parallel 1 \parallel XER_{so}$ else $(CR[CR0]) \leftarrow {}^20 \parallel 0 \parallel XER_{so}$.	9-166	9-171
stwu	RS, D(RA)	Store word (RS) in memory at EA = (RA 0) + EXTS(D). Update the base address, (RA) \leftarrow EA.	9-166	9-173
stwux	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB). Update the base address, (RA) \leftarrow EA.	9-166	9-174
stwx	RS, RA, RB	Store word (RS) in memory at EA = (RA 0) + (RB).	9-166	9-175

B.6 Arithmetic and Logical Instructions

Table B-7 lists the arithmetic and logical instructions. Arithmetic operations are performed on integer or ordinal operands stored in registers. Instructions using two operands are defined in a three-operand format, where the operation is performed on the operands stored in two registers, and the result is placed in a third register. Instructions using one operand are defined in a two-operand format, where the operation is performed on the operand in one register, and the result is placed in another register. Several instructions have immediate formats, in which one operand is coded as part of the instruction itself. Most arithmetic and logical instructions can optionally set the Condition Register (CR) based on the outcome of the instruction.

Table B-7. Arithmetic and Logical Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
add	RT, RA, RB	Add (RA) to (RB). Place result in RT.		9-6
add.			CR[CR0]	
addo			XER[SO, OV]	
addo.			CR[CR0] XER[SO, OV]	
addc	RT, RA, RB	Add (RA) to (RB). Place result in RT. Place carry-out in XER[CA].		9-7
addc.			CR[CR0]	
addco			XER[SO, OV]	
addco.			CR[CR0] XER[SO, OV]	
adde	RT, RA, RB	Add XER[CA], (RA), (RB). Place result in RT. Place carry-out in XER[CA].		9-9
adde.			CR[CR0]	
addeo			XER[SO, OV]	
addeo.			CR[CR0] XER[SO, OV]	
addi	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT.		9-9
addic	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].		9-10
addic.	RT, RA, IM	Add EXTS(IM) to (RA 0). Place result in RT. Place carry-out in XER[CA].	CR[CR0]	9-11
addis	RT, RA, IM	Add (IM ¹⁶ 0) to (RA 0). Place result in RT.		9-12

Table B-7. Arithmetic and Logical Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
addme	RT, RA	Add XER[CA], (RA), (-1). Place result in RT. Place carry-out in XER[CA].		9-13
addme.			CR[CR0]	
addmeo			XER[SO, OV]	
addmeo.			CR[CR0] XER[SO, OV]	
addze	RT, RA	Add XER[CA] to (RA). Place result in RT. Place carry-out in XER[CA].		9-14
addze.			CR[CR0]	
addzeo			XER[SO, OV]	
addzeo.			CR[CR0] XER[SO, OV]	
and	RA, RS, RB	AND (RS) with (RB). Place result in RA.		9-15
and.			CR[CR0]	
andc	RA, RS, RB	AND (RS) with \neg (RB). Place result in RA.		9-16
andc.			CR[CR0]	
andi.	RA, RS, IM	AND (RS) with ($^{16}0 \parallel$ IM). Place result in RA.	CR[CR0]	9-17
andis.	RA, RS, IM	AND (RS) with (IM \parallel $^{16}0$). Place result in RA.	CR[CR0]	9-18
cntlzw	RA, RS	Count leading zeros in RS. Place result in RA.		9-38
cntlzw.			CR[CR0]	
divw	RT, RA, RB	Divide (RA) by (RB), signed. Place result in RT.		9-59
divw.			CR[CR0]	
divwo			XER[SO, OV]	
divwo.			CR[CR0] XER[SO, OV]	
divwu	RT, RA, RB	Divide (RA) by (RB), unsigned. Place result in RT.		9-60
divwu.			CR[CR0]	
divwuo			XER[SO, OV]	
divwuo.			CR[CR0] XER[SO, OV]	
eqv	RA, RS, RB	Equivalence of (RS) with $\overline{(RB)}$. (RA) $\leftarrow \neg((RS) \oplus (RB))$		9-62
eqv.			CR[CR0]	
extsb	RA, RS	Extend the sign of byte (RS) _{24:31} . Place the result in RA.		9-63
extsb.			CR[CR0]	

Table B-7. Arithmetic and Logical Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
extsh	RA, RS	Extend the sign of halfword (RS) _{16:31} . Place the result in RA.		9-64
extsh.			CR[CR0]	
mulhw	RT, RA, RB	Multiply (RA) and (RB), signed. Place hi-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). (RT) $\leftarrow prod_{0:31}$.		9-127
mulhw.			CR[CR0]	
mulhwu	RT, RA, RB	Multiply (RA) and (RB), unsigned. Place hi-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (unsigned). (RT) $\leftarrow prod_{0:31}$.		9-128
mulhwu.			CR[CR0]	
mulli	RT, RA, IM	Multiply (RA) and IM, signed. Place lo-order result in RT. $prod_{0:47} \leftarrow (RA) \times IM$ (signed) (RT) $\leftarrow prod_{16:47}$		9-129
mullw	RT, RA, RB	Multiply (RA) and (RB), signed. Place lo-order result in RT. $prod_{0:63} \leftarrow (RA) \times (RB)$ (signed). (RT) $\leftarrow prod_{32:63}$.		9-130
mullw.			CR[CR0]	
mullwo			XER[SO, OV]	
mullwo.			CR[CR0] XER[SO, OV]	
nand	RA, RS, RB	NAND (RS) with (RB). Place result in RA.		9-131
nand.			CR[CR0]	
neg	RT, RA	Negative (two's complement) of RA. (RT) $\leftarrow \neg(RA) + 1$		9-132
neg.			CR[CR0]	
nego			XER[SO, OV]	
nego.			CR[CR0] XER[SO, OV]	
nor	RA, RS, RB	NOR (RS) with (RB). Place result in RA.		9-139
nor.			CR[CR0]	
or	RA, RS, RB	OR (RS) with (RB). Place result in RA.		9-134
or.			CR[CR0]	
orc	RA, RS, RB	OR (RS) with $\neg(RB)$. Place result in RA.		9-134
orc.			CR[CR0]	
ori	RA, RS, IM	OR (RS) with (${}^{16}0 \parallel IM$). Place result in RA.		9-142
oris	RA, RS, IM	OR (RS) with ($IM \parallel {}^{16}0$). Place result in RA.		9-143

Table B-7. Arithmetic and Logical Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
subf	RT, RA, RB	Subtract (RA) from (RB). (RT) $\leftarrow \neg(\text{RA}) + (\text{RB}) + 1$.		9-176
subf.			CR[CR0]	
subfo			XER[SO, OV]	
subfo.			CR[CR0] XER[SO, OV]	
subfc	RT, RA, RB	Subtract (RA) from (RB). (RT) $\leftarrow \neg(\text{RA}) + (\text{RB}) + 1$. Place carry-out in XER[CA].		9-177
subfc.			CR[CR0]	
subfco			XER[SO, OV]	
subfco.			CR[CR0] XER[SO, OV]	
subfe	RT, RA, RB	Subtract (RA) from (RB) with carry-in. (RT) $\leftarrow \neg(\text{RA}) + (\text{RB}) + \text{XER}[\text{CA}]$. Place carry-out in XER[CA].		9-178
subfe.			CR[CR0]	
subfeo			XER[SO, OV]	
subfeo.			CR[CR0] XER[SO, OV]	
subfic	RT, RA, IM	Subtract (RA) from EXTS(IM). (RT) $\leftarrow \neg(\text{RA}) + \text{EXTS}(\text{IM}) + 1$. Place carry-out in XER[CA].		9-179
subfme	RT, RA, RB	Subtract (RA) from (-1) with carry-in. (RT) $\leftarrow \neg(\text{RA}) + (-1) + \text{XER}[\text{CA}]$. Place carry-out in XER[CA].		9-180
subfme.			CR[CR0]	
subfmeo			XER[SO, OV]	
subfmeo.			CR[CR0] XER[SO, OV]	
subfze	RT, RA, RB	Subtract (RA) from zero with carry-in. (RT) $\leftarrow \neg(\text{RA}) + \text{XER}[\text{CA}]$. Place carry-out in XER[CA].		9-180
subfze.			CR[CR0]	
subfzeo			XER[SO, OV]	
subfzeo.			CR[CR0] XER[SO, OV]	
xor	RA, RS, RB	XOR (RS) with (RB). Place result in RA.		9-198
xor.			CR[CR0]	
xori	RA, RS, IM	XOR (RS) with ($^{16}0 \parallel \text{IM}$). Place result in RA.		9-199
xoris	RA, RS, IM	XOR (RS) with ($\text{IM} \parallel ^{16}0$). Place result in RA.		9-200

B.7 Condition Register Logical Instructions

CR logical instructions combine the results of several comparisons without incurring the overhead of conditional branching. These instructions can significantly improve code performance if multiple conditions are tested before making a branch decision. Table B-8 summarizes the CR logical instructions.

Table B-8. Condition Register Logical Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
crand	BT, BA, BB	AND bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-39
crandc	BT, BA, BB	AND bit (CR_{BA}) with $\neg(CR_{BB})$. Place result in CR_{BT} .		9-40
creqv	BT, BA, BB	Equivalence of bit CR_{BA} with CR_{BB} . $CR_{BT} \leftarrow \neg(CR_{BA} \oplus CR_{BB})$		9-41
crnand	BT, BA, BB	NAND bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-42
crnor	BT, BA, BB	NOR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-43
cror	BT, BA, BB	OR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-44
crorc	BT, BA, BB	OR bit (CR_{BA}) with $\neg(CR_{BB})$. Place result in CR_{BT} .		9-45
crxor	BT, BA, BB	XOR bit (CR_{BA}) with (CR_{BB}). Place result in CR_{BT} .		9-46
mcrf	BF, BFA	Move CR field, $(CR[CRn]) \leftarrow (CR[CRm])$ where $m \leftarrow BFA$ and $n \leftarrow BF$.		9-107

B.8 Branch Instructions

The architecture provides conditional and unconditional branches to any storage location. The conditional branch instructions test condition codes set previously and branch accordingly. Conditional branch instructions may decrement and test the Count Register (CTR) as part of determination of the branch condition and may save the return address in the Link Register (LR). The target address for a branch may be a displacement from the current instruction address (CIA), or may be contained in the LR or CTR, or may be an absolute address.

Table B-9. Branch Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
b	target	Branch unconditional relative. $LI \leftarrow (target - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^20)$		9-19
ba		Branch unconditional absolute. $LI \leftarrow target_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^20)$		
bl		Branch unconditional relative. $LI \leftarrow (target - CIA)_{6:29}$ $NIA \leftarrow CIA + EXTS(LI \parallel ^20)$	(LR) $\leftarrow CIA + 4.$	
bla		Branch unconditional absolute. $LI \leftarrow target_{6:29}$ $NIA \leftarrow EXTS(LI \parallel ^20)$	(LR) $\leftarrow CIA + 4.$	
bc	BO, BI, target	Branch conditional relative. $BD \leftarrow (target - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$	9-20
bca		Branch conditional absolute. $BD \leftarrow target_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$	
bcl		Branch conditional relative. $BD \leftarrow (target - CIA)_{16:29}$ $NIA \leftarrow CIA + EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$ (LR) $\leftarrow CIA + 4.$	
bcla		Branch conditional absolute. $BD \leftarrow target_{16:29}$ $NIA \leftarrow EXTS(BD \parallel ^20)$	CTR if $BO_2 = 0.$ (LR) $\leftarrow CIA + 4.$	
bcctr	BO, BI	Branch conditional to address in CTR. Using (CTR) at exit from instruction, $NIA \leftarrow CTR_{0:29} \parallel ^20.$	CTR if $BO_2 = 0.$	9-26
bcctrl			CTR if $BO_2 = 0.$ (LR) $\leftarrow CIA + 4.$	
bclr	BO, BI	Branch conditional to address in LR. Using (LR) at entry to instruction, $NIA \leftarrow LR_{0:29} \parallel ^20.$	CTR if $BO_2 = 0.$	9-30
bclrl			CTR if $BO_2 = 0.$ (LR) $\leftarrow CIA + 4.$	

B.9 Comparison Instructions

Comparison instructions perform arithmetic and logical comparisons between two operands and set one of the eight condition code register fields based on the outcome of the comparison. Table B-10 shows the comparison instructions supported by the PPC405 core.

Table B-10. Comparison Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
cmp	BF, 0, RA, RB	Compare (RA) to (RB), signed. Results in CR[CRn], where $n = BF$.		9-34
cmpi	BF, 0, RA, IM	Compare (RA) to EXTS(IM), signed. Results in CR[CRn], where $n = BF$.		9-35
cmpl	BF, 0, RA, RB	Compare (RA) to (RB), unsigned. Results in CR[CRn], where $n = BF$.		9-36
cmpli	BF, 0, RA, IM	Compare (RA) to ($1^{60} IM$), unsigned. Results in CR[CRn], where $n = BF$.		9-37

B.10 Rotate and Shift Instructions

Rotate and shift instructions rotate and shift operands which are stored in the general purpose registers. Rotate instructions can also mask rotated operands. Table B-11 shows the PPC405 rotate and shift instructions.

Table B-11. Rotate and Shift Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
rlwimi	RA, RS, SH, MB, ME	Rotate left word immediate, then insert according to mask. $r \leftarrow \text{ROTL}((RS), SH)$ $m \leftarrow \text{MASK}(MB, ME)$ $(RA) \leftarrow (r \wedge m) \vee ((RA) \wedge \neg m)$		9-146
rlwimi.			CR[CR0]	
rlwinm	RA, RS, SH, MB, ME	Rotate left word immediate, then AND with mask. $r \leftarrow \text{ROTL}((RS), SH)$ $m \leftarrow \text{MASK}(MB, ME)$ $(RA) \leftarrow (r \wedge m)$		9-147
rlwinm.			CR[CR0]	
rlwnm	RA, RS, RB, MB, ME	Rotate left word, then AND with mask. $r \leftarrow \text{ROTL}((RS), (RB)_{27:31})$ $m \leftarrow \text{MASK}(MB, ME)$ $(RA) \leftarrow (r \wedge m)$		9-150
rlwnm.			CR[CR0]	
slw	RA, RS, RB	Shift left (RS) by (RB) _{27:31} . $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), n)$. if (RB) ₂₆ = 0 then $m \leftarrow \text{MASK}(0, 31 - n)$ else $m \leftarrow {}^{32}0$. $(RA) \leftarrow r \wedge m$.		9-152
slw.			CR[CR0]	
sraw	RA, RS, RB	Shift right algebraic (RS) by (RB) _{27:31} . $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. if (RB) ₂₆ = 0 then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow {}^{32}0$. $s \leftarrow (RS)_0$. $(RA) \leftarrow (r \wedge m) \vee ({}^{32}s \wedge \neg m)$. $\text{XER}[CA] \leftarrow s \wedge ((r \wedge \neg m) \neq 0)$.		9-153
sraw.			CR[CR0]	
srawi	RA, RS, SH	Shift right algebraic (RS) by SH. $n \leftarrow SH$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. $m \leftarrow \text{MASK}(n, 31)$. $s \leftarrow (RS)_0$. $(RA) \leftarrow (r \wedge m) \vee ({}^{32}s \wedge \neg m)$. $\text{XER}[CA] \leftarrow s \wedge ((r \wedge \neg m) \neq 0)$.		9-154
srawi.			CR[CR0]	
srw	RA, RS, RB	Shift right (RS) by (RB) _{27:31} . $n \leftarrow (RB)_{27:31}$. $r \leftarrow \text{ROTL}((RS), 32 - n)$. if (RB) ₂₆ = 0 then $m \leftarrow \text{MASK}(n, 31)$ else $m \leftarrow {}^{32}0$. $(RA) \leftarrow r \wedge m$.		9-155
srw.			CR[CR0]	

B.11 Cache Control Instructions

Cache control instructions allow the user to indirectly control the contents of the data and instruction caches. The user may fill, flush, invalidate and zero blocks (16-byte lines) in the data cache. The user may also invalidate congruence classes in both caches and invalidate individual lines in the instruction cache.

Table B-12. Cache Control Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
dcba	RA, RB	Speculatively establish the data cache block which contains the EA (RA 0) + (RB).		9-47
dcbf	RA, RB	Flush (store, then invalidate) the data cache block which contains the EA (RA 0) + (RB).		9-49
dcbi	RA, RB	Invalidate the data cache block which contains the EA (RA 0) + (RB).		9-50
dcbst	RA, RB	Store the data cache block which contains the EA (RA 0) + (RB).		9-51
dcbt	RA, RB	Load the data cache block which contains the EA (RA 0) + (RB).		9-52
dcbtst	RA, RB	Load the data cache block which contains the EA (RA 0) + (RB).		9-53
dcbz	RA, RB	Zero the data cache block which contains the EA (RA 0) + (RB).		9-54
dccci	RA, RB	Invalidate the data cache congruence class associated with the EA (RA 0) + (RB).		9-56
dcread	RT, RA, RB	Read either tag or data information from the data cache congruence class associated with the EA (RA 0) + (RB). Place the results in RT.		9-57
icbi	RA, RB	Invalidate the instruction cache block which contains the EA (RA 0) + (RB).		9-65
icbt	RA, RB	Load the instruction cache block which contains the EA (RA 0) + (RB).		9-66
iccci	RA, RB	Invalidate instruction cache.		9-67
icread	RA, RB	Read either tag or data information from the instruction cache congruence class associated with the EA (RA 0) + (RB). Place the results in ICDBDR.		9-68

B.12 Interrupt Control Instructions

The interrupt control instructions allow the user to move data between general purpose registers and the machine state register, return from interrupts and enable or disable maskable external interrupts. Table B-13 shows the interrupt control instruction set.

Table B-13. Interrupt Control Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
mfmsr	RT	Move from MSR to RT, (RT) ← (MSR).		9-111
mtmsr	RS	Move to MSR from RS, (MSR) ← (RS).		9-118
rfci		Return from critical interrupt (PC) ← (SRR2). (MSR) ← (SRR3).		9-144
rfi		Return from interrupt. (PC) ← (SRR0). (MSR) ← (SRR1).		9-144
wrtee	RS	Write value of RS ₁₆ to the External Enable bit (MSR[EE]).		9-196
wrteei	E	Write value of E to the External Enable bit (MSR[EE]).		9-197

B.13 TLB Management Instructions

The TLB management instructions read and write entries of the TLB array in the MMU, search the TLB array for an entry which will translate a given address, invalidate all TLB entries, and synchronize TLB updates with other processors.

Table B-14. TLB Management Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
tlbia		All of the entries in the TLB are invalidated and become unavailable for translation by clearing the valid (V) bit in the TLBHI portion of each TLB entry. The rest of the fields in the TLB entries are unmodified.		9-183

Table B-14. TLB Management Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
tlbre	RT, RA, WS	<p>If WS = 0: Load TLBHI portion of the selected TLB entry into RT. Load the PID register with the contents of the TID field of the selected TLB entry. $(RT) \leftarrow TLBHI[(RA)]$ $(PID) \leftarrow TLB[(RA)]_{TID}$</p> <p>If WS = 1: Load TLBLO portion of the selected TLB entry into RT. $(RT) \leftarrow TLBLO[(RA)]$</p>		9-184
tlbsx	RT, RA, RB	<p>Search the TLB array for a valid entry which translates the EA $EA = (RA 0) + (RB)$. If found, $(RT) \leftarrow$ Index of TLB entry. If not found, (RT) Undefined.</p>		9-186
tlbsx.		<p>If found, $(RT) \leftarrow$ Index of TLB entry. $CR[CR0]_{EQ} \leftarrow 1$. If not found, (RT) Undefined. $CR[CR0]_{EQ} \leftarrow 1$.</p>	CR[CR0] _{LT,GT,SO}	
tlbsync		<p>tlbsync does not complete until all previous TLB-update instructions executed by this processor have been received and completed by all other processors.</p> <p>For the PPC405 core, tlbsync is a no-op.</p>		9-187
tlbwe	RS, RA, WS	<p>If WS = 0: Write TLBHI portion of the selected TLB entry from RS. Write the TID field of the selected TLB entry from the PID register. $TLBHI[(RA)] \leftarrow (RS)$ $TLB[(RA)]_{TID} \leftarrow (PID)_{24:31}$</p> <p>If WS = 1: Write TLBLO portion of the selected TLB entry from RS. $TLBLO[(RA)] \leftarrow (RS)$</p>		9-188

B.14 Processor Management Instructions

The processor management instructions move data between GPRs and SPRs and DCRs in the PPC405 core; these instructions also provide traps, system calls and synchronization controls.

Table B-15. Processor Management Instructions

Mnemonic	Operands	Function	Other Registers Changed	Page
eieio		Storage synchronization. All loads and stores that precede the eieio instruction complete before any loads and stores that follow the instruction access main storage. Implemented as sync , which is more restrictive.		9-61
isync		Synchronize execution context by flushing the prefetch queue.		9-70
mcrxr	BF	Move XER[0:3] into field CR _n , where $n \leftarrow BF$. $CR[CR_n] \leftarrow (XER[SO, OV, CA])$. $(XER[SO, OV, CA]) \leftarrow ^30$.		9-108
mfcrr	RT	Move from CR to RT, $(RT) \leftarrow (CR)$.		9-108
mfdcr	RT, DCRN	Move from DCR to RT, $(RT) \leftarrow (DCR(DCRN))$.		9-110
mfspr	RT, SPRN	Move from SPR to RT, $(RT) \leftarrow (SPR(SPRN))$.		9-112
mtcrf	FXM, RS	Move some or all of the contents of RS into CR as specified by FXM field, $mask \leftarrow ^4(FXM_0) \parallel ^4(FXM_1) \parallel \dots \parallel ^4(FXM_6) \parallel ^4(FXM_7)$. $(CR) \leftarrow ((RS) \wedge mask) \vee (CR) \wedge \neg mask$.		9-116
mtdcr	DCRN, RS	Move to DCR from RS, $(DCR(DCRN)) \leftarrow (RS)$.		9-117
mtspr	SPRN, RS	Move to SPR from RS, $(SPR(SPRN)) \leftarrow (RS)$.		9-119
sc		System call exception is generated. $(SRR1) \leftarrow (MSR)$ $(SRR0) \leftarrow (PC)$ $PC \leftarrow EVPR_{0:15} \parallel 0x0C00$ $(MSR[WE, PR, EE, PE, DR, IR]) \leftarrow 0$		9-151
sync		Synchronization. All instructions that precede sync complete before any instructions that follow sync begin. When sync completes, all storage accesses initiated before sync will have completed.		9-182
tw	TO, RA, RB	Trap exception is generated if, comparing (RA) with (RB), any condition specified by TO is true.		9-190

Table B-15. Processor Management Instructions (continued)

Mnemonic	Operands	Function	Other Registers Changed	Page
twi	TO, RA, IM	Trap exception is generated if, comparing (RA) with EXTS(IM), any condition specified by TO is true.		9-193

Appendix C. Code Optimization and Instruction Timings

The code optimization guidelines in “Code Optimization Guidelines” and the information describing instruction timings in “Instruction Timings,” on page C-3 can help compiler, system, and application programmers produce high-performance code and determine accurate execution times.

C.1 Code Optimization Guidelines

The following guidelines can help to reduce program execution times.

C.1.1 Condition Register Bits for Boolean Variables

Compilers can use Condition Register (CR) bits to store boolean variables, where 0 and 1 represent False and True values, respectively. This generally improves performance, compared to using General Purpose Registers (GPRs) to store boolean variables. Most common operations on boolean variables can be accomplished using the CR Logical instructions.

C.1.2 CR Logical Instruction for Compound Branches

For example, consider the following pseudocode:

```
if (Var28 || Var29 || Var30 || Var 31) branch to target
```

Var28–Var31 are boolean variables, maintained as bits in the CR[CR7] field (CR_{28:31}). The value 1 represents True; 0 represents False.

This could be coded with branches as:

```
bt      28, target
bt      29, target
bt      30, target
bt      31, target
```

Generally faster, functionally equivalent code, using CR Logical instructions, follows:

```
crcr    2, 28, 29
cror    2, 2, 30
cror    2, 2, 31
bt      2, target
```

C.1.3 Floating-Point Emulation

Two ways of handling floating-point emulation are available.

The preferred method is a call interface to subroutines in a floating-point emulation run-time library.

Alternatively, code can use the PowerPC floating point instructions. The PPC405, an integer processor, does not recognize these instructions and will take an illegal instruction interrupt. The interrupt handler can be written to determine the instruction opcode and execute appropriate (integer-based) library routines to provide the equivalent function.

Because this method adds interrupt context switching time to the execution time of library routines that would have been called directly by the preferred method, it is not preferred. However, this method supports code that contains PowerPC floating-point instructions.

C.1.4 Cache Usage

Code and data can be organized, based on the size and structure of the instruction and data cache arrays, to minimize cache misses.

In the cache arrays, any two addresses in which $A_{m:26}$ (the index) are the same, but which differ in $A_{0:m-1}$ (the tag), are called congruent. (This describes a two-way set-associative cache.) $A_{27:31}$ define the 32 bytes in a cache line, the smallest object that can be brought into the cache. Only two congruent lines can be in the cache simultaneously; accessing a third congruent line causes the removal from the cache of one of the two lines previously there

Table C-1 illustrates the value of m and the index size for the various cache array sizes.

Table C-1. Cache Sizes, Tag Fields, and Lines

Array Size	Instruction Cache Array			Data Cache Array		
	m (Tag Field Bits)	n (Lines)	Index Bits	m (Tag Field Bits)	n (Lines)	Index Bits
0KB	—	—	—	—	—	—
4KB	22 (0:21)	64	21:26	20 (0:19)	64	21:26
8KB	22 (0:21)	128	20:26	20 (0:19)	128	20:26
16KB	22 (0:21)	256	19:26	20 (0:19)	256	19:26
32KB	22 (0:21)	512	18:26	20 (0:19)	512	18:26

Moving new code and data into the cache arrays occurs at the speed of external memory. Much faster execution is possible when all code and data is available in the cache. Organizing code to uniformly use $A_{m:26}$ minimizes the use of congruent addresses.

C.1.5 CR Dependencies

For CR-setting arithmetic, compare, CR-logical, and logical instructions, and the CR-setting **mcrf**, **mcrxr**, and **mtcrf** instructions, put two instructions between the CR-setting instruction and a Branch instruction that uses a bit in the CR field set by the CR-setting instruction.

C.1.6 Branch Prediction

Use the Y-bit in branch instructions to force proper branch prediction when there is a more likely prediction than the standard prediction. See “Branch Prediction” on page 2-26 for a more information about branch prediction.

C.1.7 Alignment

For speed, align all accesses on the appropriate operand-size boundary. For example, load/store word operands should be word-aligned, and so on. Hardware does not trap unaligned accesses; instead, two accesses are performed for a load or store of an unaligned operand that crosses a word boundary. Unaligned accesses that do not cross word boundaries are performed in one access.

Align branch targets that are unlikely to be hit by “fall-through” code on cache line boundaries (such as the address of functions such as **strcpy**), to minimize the number of unused instructions in cache line fills.

C.2 Instruction Timings

The following timing descriptions consider only “first order” effects of cache misses in the ICU (instruction-side) and DCU (data-side) arrays.

The timing descriptions *do not* provide complete descriptions of the performance penalty associated with cache misses; the timing descriptions do not consider bus contention between the instruction-side and the data-side, or the time associated with performing line fills or flushes. Unless specifically stated otherwise, the number of cycles apply to systems having zero-wait memory access.

C.2.1 General Rules

Instructions execute in order.

All instructions, assuming cache hits, execute in one cycle, except:

- Divide instructions execute in 35 clock cycles.
- Branches execute in one or three clock cycles, as described in “Branches.”
- MAC and multiply instructions execute in one to five cycles as described in “Multiplies.”
- Aligned load/store instructions that hit in the cache execute in one clock cycle/word. See “Alignment” for information on execution timings for unaligned load/stores.
- In isolation, a data cache control instruction takes two cycles in the processor pipeline. However, subsequent DCU accesses are stalled until a cache control instruction finishes accessing the data cache array.

Note: Note that subsequent DCU accesses do not remain stalled while transfers associated with previous data cache control instructions continue on the PLB.

C.2.2 Branches

Branch instructions are decoded in prefetch buffer 0 (PFB0) and the decode stage of the instruction pipeline. Branch targets, whether the branch is known or predicted taken, can be fetched from the PFB0 and DCD stages. Incorrectly predicted branches can be corrected from the DCD or EXE (execute) stages of the pipeline.

Branches can be known taken or known not taken, or can have address or condition dependencies. Branches having address dependencies are never predicted taken. The directions of conditional branches having no address dependencies are statically predicted.

Conditional branches may depend on the results of an instruction that is changing the CR or the CTR.

Address dependencies can occur when:

- A **bclr** instruction that is known taken, or unresolved, follows (immediately, or separated by only one instruction) a link updating instruction (**mtlr** or a branch and link).
- A **bcctr** instruction that is known taken, or unresolved, follows (immediately, or separated by only one instruction) a counter updating instruction (**mtctr** or a branch that decrements the counter).

Instruction timings for branch instructions follow:

- A branch known not taken (BKNT) executes in one clock cycle. By definition a BKNT does not have address or condition dependencies.
- A branch known taken (BKT) by definition has no condition dependencies, but can have address dependencies. A BKT without address dependencies can execute in one clock cycle if it is first decoded from the PFB0 stage, or in two clock cycles if it is first decoded in the DCD stage. A BKT having address dependencies can execute in two clock cycles if there is one instruction between the branch and the address dependency, or in three clock cycles if there are no instructions between the branch and address dependency.
- A branch predicted not taken (BPNT), which must have condition dependencies, executes in one clock cycle if the prediction is correct. If the prediction is incorrect, the branch can take two or three cycles. If there was one instruction between the branch and the instruction causing the condition dependency, the branch executes in two cycles. If there were no instructions between the branch and the instruction causing the condition dependency, the branch executes in three clock cycles.
- A branch that is correctly predicted taken (BPT), which must have condition dependencies, executes in one clock cycle, if it is first decoded from the PFB0 stage, or two clock cycles if it is first decoded in the DCD stage. If the prediction is incorrect, the branch can take two or three cycles. If there is one instruction between the branch and the instruction causing the condition dependency, the branch executes in two cycles. If there are no instructions between the branch and the instruction causing the condition dependency, the branch executes in three clock cycles.

C.2.3 Multiplies

For multiply instructions having two word operands, hardware internal to the core automatically detects smaller operand sizes (by examining sign bit extension) to reduce the number of cycles necessary to complete the multiplication.

The PPC405 also supports multiply accumulate (MAC) instructions and multiply instructions having halfword operands.

Word and halfword multiply instructions are pipelined in the execution unit and use the same multiplication hardware. Because these instructions are pipelined in the execution stage they have latency and reissue rate cycle numbers. Under conditions to be described, a second multiply or MAC instruction can begin execution before the first multiply or MAC instruction completes. When these conditions are met, the reissue rate cycle numbers should be used; otherwise, the latency cycle numbers should be used. (A MAC or multiply instruction can follow another MAC or a multiply and still meet the conditions that support the use of the reissue rate cycle numbers.

Use reissue rate cycle numbers for multiply or MAC instructions that are followed by another multiply or MAC instruction, and do not have an operand dependency from a previous multiply or MAC instruction. However, one operand dependency is allowed for reissue rate cycle numbers. Internal forwarding logic allows the accumulate value of a first MAC instruction to be used as the accumulate value of a second MAC instruction without affecting the reissue rate.

Use latency cycle numbers for multiply or MAC instructions that are not followed by another multiply or MAC, or that have an operand dependency from a previous multiply or MAC instruction. However, accumulate-only dependencies between adjacent MAC instructions use reissue rate cycle numbers.

An operand dependency exists when a second multiply or MAC instruction depends on the result of a first multiply or MAC instruction.

Table C-2 summarizes the multiply and MAC instruction timings. In the table, the syntax “[o]” indicates that the instruction has an “o” form that updates XER[SO,OV], and a “non-o” form. The syntax “[.]” indicates that the instruction has a “record” form that updates CR[CR0], and a “non-record” form.

Table C-2. Multiply and MAC Instruction Timing

Operation	Reissue Rate Cycles	Latency Cycles
MAC		
MAC and negative MAC instructions	1	2
Halfword × Halfword		
mullhw[.], mullhwu[.], mulhww[.], mulhwwu[.], mulchw[.], mulchwu[.]	1	2
mulli[.], mullw[o][.], mulhw[.], mulhwu[.]	2	3
Halfword × Word		
mulli[.], mullw[o][.], mulhw[.], mulhwu[.]	2	3
Word × Word		
mullw[o][.], mulhw[.], mulhwu[.]	4	5

C.2.4 Scalar Load Instructions

Generally, the PPC405 executes cachable load instructions that hit in the data cache array or line fill buffer, or noncachable load instructions that hit in the line fill buffer (when enabled), in one cycle. However, the pipelined nature of load instructions can even cause loads that hit in the cache or line fill buffer to appear to take extra cycles under some conditions.

If a load is followed by an instruction that uses the load target as an operand, a load-use dependency exists. When the load target is returned, it is forwarded to the operand register of the “using” instruction. This forwarding results in an additional cycle of latency to a load immediately followed by a “using” instruction, causing the load to appear to execute in two cycles.

To improve cache-to-core timing or data-side on-chip memory (OCM)- to-core timing, the system designer can disable operand forwarding from the data cache unit (DCU) or OCM to the core. When operand forwarding is disabled, the load data needed by the “using” instruction is placed in an intermediate latch before the load data is forwarded to the operand register of the “using” instruction. When the load target is returned, it is forwarded to the operand register of the “using” instruction. This introduces two additional cycles of latency to a load immediately followed by a “using” instruction, causing the load instruction to appear to execute in three cycles.

Because the PPC405 can execute instructions that follow load misses if no load-use dependency exists, the load and the “using” instruction should be separated by two “non-using” instructions when possible. If only one instruction can be placed between the load and the “using” instruction, the load appears to execute in two cycles.

C.2.5 Scalar Store Instructions

Cachable stores that miss in the DCU, and noncachable stores, are queued in the data cache so that the store appears to execute in a single cycle if operand-aligned. Under certain conditions, the DCU can pipeline up to three store instructions. (See Chapter 4, “Cache Operations,” for more information.) **stwcx.** instructions that do not cause alignment errors execute in two cycles.

C.2.6 Alignment in Scalar Load and Store Instructions


The PPC405 requires an extra cycle to execute scalar loads and stores having unaligned big or little endian data (except for **lwarx** and **stwcx.**, which require word-aligned operands). If the target data is not operand aligned, and the sum of the least two significant bits of the effective address (EA) and the byte count is greater than four, the PPC405 decomposes a load or store scalar into two load or store operations. That is, the PPC405 never presents the DCU with a request for a transfer that crosses a word boundary. For example, a **lwz** with an EA of 0b11 causes the PPC405 to decompose the **lwz** into two load operations. The first load operation is for a byte at the starting effective address; the second load operation is for three bytes, starting at the next word address.

C.2.7 String and Multiple Instructions

Calculating execution times for string and multiple instructions (**lmw** and **stmw**) instructions requires an understanding of data alignment, and of the behavior of the string instructions with respect to alignment.

In the following example, the string contains 21 bytes. The first three bytes do not begin on a word boundary, and the final two bytes do not end on a word boundary. The PPC405 handles any unaligned leading bytes as a special case, then moves as many bytes as aligned words as possible, and finally handles any unaligned trailing bytes as a special case.

In the following example, arrows indicate word boundaries (the address is an exact multiple of four); shaded boxes represent unaligned bytes.



The execution time of the string instruction is the sum of the:

1. Cycles required to handle unaligned leading bytes; if any, add one clock cycle.

In the example, there are unaligned leading bytes; this transfer adds one clock cycle.

2. Cycles required to handle the number of word-aligned transfers required. Assuming data cache hits, each word-aligned transfer requires one clock cycle.

In the example, there are four aligned words; this transfer requires four clock cycles.

3. Cycles required to handle unaligned trailing bytes; if any, add one clock cycle.

In the example, there are unaligned trailing bytes; this transfer adds one clock cycle.

A string instruction operating on the example 21-byte string requires six clock cycles.

C.2.8 Loads and Store Misses

Cachable stores that miss in the DCU, and noncachable stores, are queued internally in the DCU so that the store instruction appears to execute in one cycle. Under certain conditions, the DCU can pipeline up to three store instructions. (See the Chapter 4, “Cache Operations,” for more information.)

Because the PPC405 can execute instructions that follow load misses if no load-use dependency exists, the load and the “using” instruction should be separated by “non-using” instructions whenever possible. The number of load miss penalty cycles incurred by a load that misses in the DCU or DCU line fill buffer is reduced by one cycle for every non-use instruction following the load. When the number of non-use instructions following the load is equal to or greater than the number of cycles that it takes to obtain the load data, the load instruction appears to execute in a single cycle. The number of cycles that it takes to obtain load data when it misses in the data cache and line fill buffer depends on whether operand forwarding is enabled or disabled and the system memory timing.

C.2.9 Instruction Cache Misses

Refer to “Instruction Processing” on page 2-23 for detailed information about the instruction queue and instruction fetching. Table C-3 illustrates instruction cache penalties for cachable and noncachable fetches that miss in the ICU array and line fill buffer.

Table C-3. Instruction Cache Miss Penalties

Type of ICU Request	Miss Penalty Cycles
Sequential	3
Branch Taken from DCD	5
Branch Taken from PFB0	4

Table C-3 assumes that:

- The PPC405 and processor local bus (PLB) run at the same frequency
- The PLB returns an address acknowledge during the first cycle in which the DCU asserts the PLB request
- The target instruction is returned in the cycle following the address acknowledge cycle

The penalty cycles shown for sequential ICU requests assume that the DCD stage and pre-fetch queue are filled with single-cycle nonbranching instructions or BKNT branch instructions. The penalty cycles for the remaining two rows are for taken branches from DCD and PFB0, respectively.

Index

A

- AA field
 - conditional branches 2-24
 - unconditional branches 2-24
- access protection
 - cache instructions 7-16
 - string instructions 7-17
 - virtual mode 7-13
- add 9-6
- add. 9-6
- addc 9-7
- addc. 9-7
- addco 9-7
- addco. 9-7
- adde 9-8
- adde. 9-8
- addeo 9-8
- addeo. 9-8
- addi 9-9
- addic 9-10
- addic. 9-11
- addis 9-12
- addme 9-13
- addme. 9-13
- addmeo 9-13
- addmeo. 9-13
- addo 9-6
- addo. 9-6
- address translation
 - illustrated 7-2
 - MMU 7-1
 - relationship between TLBs, illustrated 7-9
- addressing modes 1-10
- addze 9-14
- addze. 9-14
- addzeo 9-14
- addzeo. 9-14
- alignment
 - for cache control instructions 2-16
 - for storage reference instructions 2-16
 - of data types 2-16
- alignment interrupts
 - causes of 2-17
 - register settings 5-19
 - summary 5-19
- and 9-15
- and. 9-15
- andc 9-16
- andc. 9-16
- andi. 9-17
- andis. 9-18
- architecture, PowerPC 1-3
- arithmetic compares 2-11
- arithmetic instructions 2-38
- asynchronous interrupts 5-1

B

- b 9-19
- ba 9-19
- bc 9-20
- bca 9-20
- bcctr 9-26
- bcctrl 9-26
- bcl 9-20
- bcla 9-20
- bclr 9-30
- bclrl 9-30
- bctr 9-27
- bctrl 9-27
- bdnz 9-21
- bdnza 9-21
- bdnzf 9-21
- bdnzfa 9-21
- bdnzfl 9-21
- bdnzfla 9-21
- bdnzflr 9-31
- bdnzflrl 9-31
- bdnzl 9-21
- bdnzla 9-21
- bdnzlr 9-31
- bdnzlrl 9-31
- bdnzt 9-21
- bdnzta 9-21
- bdnztl 9-21
- bdnztla 9-21
- bdnztlr 9-31
- bdnztlrl 9-31
- bdz 9-21
- bdza 9-21
- bdzf 9-22
- bdzfa 9-22
- bdzfl 9-22
- bdzfla 9-22
- bdzflr 9-31
- bdzflrl 9-31
- bdzl 9-21
- bdzla 9-21
- bdzlr 9-31
- bdzlrl 9-31
- bdzt 9-22
- bdzta 9-22
- bdztl 9-22
- bdztla 9-22
- bdztlr 9-31
- bdztlrl 9-31
- beq 9-22
- beqa 9-22
- beqctr 9-27
- beqctrl 9-27
- beql 9-22
- beqlr 9-31
- beqlrl 9-31
- bf 9-22

bfa 9-22
 bfctr 9-27
 bfctrl 9-27
 bfl 9-22
 bfla 9-22
 bflr 9-32
 bflrl 9-32
 bge 9-23
 bgea 9-23
 bgectrl 9-27
 bgei 9-23
 bgela 9-23
 bgeir 9-32
 bgeirl 9-32
 bgrctr 9-27
 bgt 9-23
 bgta 9-23
 bgctr 9-27
 bgctrl 9-27
 bgtl 9-23
 bgtla 9-23
 bgtlr 9-32
 bgtlrl 9-32
 BI field
 conditional branches 2-25
 big endian
 alignment 2-17
 defined 2-18
 mapping 2-19
 storage regions
 byte-reverse instructions 2-21- 2-22
 bi 9-19
 bla 9-19
 ble 9-23
 blea 9-23
 blectr 9-27
 blectrl 9-27
 blel 9-23
 blela 9-23
 bleir 9-32
 bleirl 9-32
 blr 9-30
 blrl 9-30
 blt 9-23
 blta 9-23
 bltctr 9-27
 bltctrl 9-27
 bitl 9-23
 bitla 9-23
 bitlr 9-32
 bitlrl 9-32
 bne 9-24
 bnea 9-24
 bnectr 9-28
 bnectrl 9-28
 bnel 9-24
 bnela 9-24
 bnelr 9-32
 bnelrl 9-32
 bng 9-24
 bnga 9-24
 bngctr 9-28
 bngctrl 9-28
 bngl 9-24
 bngla 9-24
 bnglr 9-32
 bnglrl 9-32
 bnl 9-24
 bnla 9-24
 bnlctr 9-28
 bnlctrl 9-28
 bnll 9-24
 bnlla 9-24
 bnllr 9-33
 bnllrl 9-33
 bns 9-24
 bnsa 9-24
 bnsctr 9-28
 bnsctrl 9-28
 bnsi 9-24
 bnsia 9-24
 bnsir 9-33
 bnsirl 9-33
 bnu 9-25
 bnua 9-25
 bnuctr 9-28
 bnuctrl 9-28
 bnul 9-25
 bnula 9-25
 bnulr 9-33
 bnulrl 9-33
 BO field
 conditional branches 2-25
 branch instructions 2-40
 branch prediction 2-26, A-1, B-9
 controlling through mnemonics 2-27
 branching control
 AA field on conditional branches 2-24
 AA field on unconditional branches 2-24
 BI field on conditional branches 2-25
 BO field on conditional branches 2-25
 branch prediction 2-26
 bso 9-25
 bsoa 9-25
 bsoctr 9-28
 bsoctrl 9-28
 bsol 9-25
 bsola 9-25
 bsolr 9-33
 bsolrl 9-33
 bt 9-25
 bta 9-25
 btctr 9-28
 btctrl 9-28
 btl 9-25
 btla 9-25
 btlr 9-33
 btlrl 9-33
 bun 9-25
 buna 9-25
 bunctr 9-29
 bunctrl 9-29

- bunl 9-25
- bunla 9-25
- bunlr 9-33
- bunlrl 9-33
- byte ordering
 - big endian, defined 2-18
 - little endian
 - defined 2-18
 - supported 2-19
 - overview 2-17
- byte reversal
 - during load/store access 2-21
- byte-reverse instructions
 - augmented by endian (E) storage attribute 2-23
 - compare to endian (E) storage attribute 2-21
- C**
- cache
 - instructions
 - DAC debug events 8-15
 - cache block, defined 4-9
 - cache control instructions
 - access protection 7-16
 - causing data storage interrupts 7-16
 - cache line
 - dirty, defined 4-16
 - See also cache block
 - cache line fills
 - DCU 4-6
 - defined 4-6
 - types 4-4
 - caches. See ICU;DCU
 - caching inhibited (I) storage attribute
 - for data accesses, controlled by DCCR 7-20
 - for instruction fetches, controlled by ICCR 7-20
 - virtual mode 7-5
- CCR0 10-6
- clrlswi 9-147
- clrlswi. 9-147
- clrlwi 9-147
- clrlwi. 9-147
- clrrwi 9-148
- clrrwi. 9-148
- cmp 9-34
- cmpi 9-35
- cmpl 9-36
- cmpli 9-37
- cmplw 9-36
- cmplwi 9-37
- cmpw 9-34
- cmpwi 9-35
- cntlzw 9-38
- cntlzw. 9-38
- compare instructions
 - arithmetic 2-11
 - in core, listed 2-39
 - effect on CR fields 2-12
 - logical 2-11
- Condition Register. See CR
- conditional branches
 - AA field 2-24
 - BI field 2-25
 - BO field 2-25
 - mnemonics used to control prediction 2-27
- context synchronization
 - defined 2-33
 - for ITLB 7-7
 - limitations 2-33
- context, defined 2-33
- conventions
 - notational xxii
- Count Register. See CTR
- CR 10-8
- CR (Condition Register)
 - arithmetic and logical instructions 2-38
 - compare instructions 2-11, 2-39
 - CR0 field 2-12– 2-13
 - logical instructions 2-39
 - setting fields 2-10
 - summarized 1-9
- crand 9-39
- crandc 9-40
- crclr 9-46
- creqv 9-41
- critical input interrupts 5-13
 - register settings 5-14
- critical interrupts 5-3
 - defined 5-5
 - processing 5-6
- crmave 9-44
- crnand 9-42
- crnor 9-43
- crnot 9-43
- cror 9-44
- crorc 9-45
- crset 9-41
- crxor 9-46
- CTR 10-9
- CTR (Count Register)
 - branch instructions 2-40
 - functions 2-6
 - testing by branch instructions 2-25
- D**
- DAC1 8-9
- DAC1–DAC2 8-9, 10-10
- Data Address Compare Register (DAC1) 8-9
- data alignment
 - in little endian storage 2-21
 - overview 2-16
- Data Cache Cachability Register. See DCCR
- Data Cache Cachability Register. See DCCR
- data cache unit. See DCU
- Data Cache Write-through Register. See DCWR
- Data Cache Write-through Register. See DCWR
- Data Exception Address Register. See DEAR
- data machine check interrupts
 - register settings 5-15
- data storage interrupts
 - caused by cache control instructions 7-16
 - causes 5-16
 - described 7-10
 - programming note 5-16
 - register settings 5-17

- data TLB. *See* DTLB
- data types
 - illustrated 2-16
 - summarized 1-8
- DBCR 8-4
- DBCR0 10-11, 10-13
- DBCR0 (Debug Control Register 0)
 - resets 3-1
- DBSR 8-7, 10-15
- dcbi
 - does not cause interrupts 7-16
 - functions 4-10
- dcbf 9-49
 - data storage interrupts 7-17
 - functions 4-10
- dcbi 9-50
 - data storage interrupts 7-16
 - functions 4-10
- dcbst 9-51
 - data storage interrupts 7-17
 - functions 4-10
- dcbt 9-52
 - data storage interrupts 7-17
 - functions 4-10
- dcbtst
 - functions 4-10
- dcbz 9-54
 - data storage interrupts 7-16
 - functions 4-11
- dccci 9-56
 - data storage interrupts 7-16
 - functions 4-11
 - when use not recommended 7-17
- DCCR 10-17
- DCCR (Data Cache Cachability Register)
 - controlling cachability 4-8
 - controlling the caching inhibited (I) storage attribute 7-20
- dcread 9-57
 - controlled by CCR0 4-11
 - as debugging tool 4-15
 - functions 4-11
- DCRs (device control registers)
 - instructions for reading 2-32
 - summary 1-9
 - uses for 2-15
- DCU (data cache unit)
 - cachability control 4-8
 - cache line fills 4-6
 - coherency 4-9
 - debugging 4-15
 - features 4-1
 - instructions 4-10
 - load commands, accepting 4-17
 - load strategies 4-8
 - overview 4-6
 - performance 4-16
 - pipeline stalls 4-16
 - priority changes 4-17
 - priority signal 4-17
 - sequential caching 4-18
 - simultaneous cache operations 4-17
 - store commands 4-17
 - tag information in GPRs 4-16
 - write strategies 4-7
- DCWR 10-19
- DCWR (Data Cache Write-through Register)
 - controlling write strategies 4-7
 - write-through policy 7-19
- DEAR 10-21
- DEAR (Data Exception Address Register)
 - illustrated 5-13
- Debug Control Register (DBCR) 8-4
- Debug Control Register 0. *See* DBCR0
- debug interrupts
 - register settings 5-26
- debugging
 - boundary scan chain 8-21
 - DCU 4-15
 - debug events 8-10
 - debug interfaces 8-19
 - JTAG test access port 8-19
 - trace status port 8-22
 - development tools 8-1
 - ICU 4-14
 - modes 8-1
 - external 8-2
 - internal 8-1
 - real-time trace 8-3
 - processor control 8-3
 - processor status 8-4
- Device Control Registers. *See* DCRs
- dirty cache line, defined 4-16
- divw 9-59
- divw. 9-59
- divwo 9-59
- divwo. 9-59
- divwu 9-60
- divwu. 9-60
- divwuo 9-60
- divwuo. 9-60
- DTLB (data translation lookaside buffer)
 - accesses 7-7
 - miss interrupts 5-25, 7-11
 - summary 7-7

E

- EA (effective address)
 - forming 2-16
 - translation to RA, illustrated 7-2
 - when non-cachable 4-9
- EAs (effective addresses)
 - indexing the cache array 4-5
- effective address. *See* EA
- effective addresses. *See* EAs
- eieio 9-61
 - storage synchronization 2-36
- embedded processors
 - instruction set 2-37
- endian (E) storage attribute
 - and byte-reverse load/store instructions 2-23
 - controlled by SLER 7-20
 - and little endian 2-19

- when controlled by TLB 7-6
- engineering note
 - ESR bits 5-13
- eqv 9-62
- eqv. 9-62
- ESR 10-23
- ESR (Exception Status Register)
 - usage for program interrupts 5-20
- ESR (Exception Syndrome Register)
 - clearing privileged exceptions 2-31
 - engineering note 5-13
 - illustrated 5-11
 - MCI bit, behavior of 5-12
- EVPR 10-25
- EVPR (Exception Vector Prefix Register)
 - illustrated 5-10
- Exception Syndrome Register. *See* ESR
- Exception Vector Prefix Register. *See* EVPR
- exceptions
 - defined 5-1
 - handling, and MSR bits 2-31
 - privileged, clearing 2-31
 - registers during debug exceptions 5-26
- exceptions. *See also* interrupts
- execution mode
 - controlling by MSR 2-31
- execution synchronization, defined 2-35
- extended memonics
 - beqlr 9-31
- extended menmonics
 - blectrl 9-27
 - bnlctrl 9-28
- extended mnemoniccd
 - bngla 9-24
- extended mnemonics
 - alphabetical B-9
 - bctr 9-27
 - bctrl 9-27
 - bdnz 9-21
 - bdnza 9-21
 - bdnzf 9-21
 - bdnzfa 9-21
 - bdnzfkr 9-31
 - bdnzfl 9-21
 - bdnzfla 9-21
 - bdnzflrl 9-31
 - bdnzl 9-21
 - bdnzla 9-21
 - bdnzlr 9-31
 - bdnzlrl 9-31
 - bdnzt 9-21
 - bdnzta 9-21
 - bdnztl 9-21
 - bdnztla 9-21
 - bdnztlr 9-31
 - bdnztlrl 9-31
 - bdz 9-21
 - bdza 9-21
 - bdzf 9-22
 - bdzfa 9-22
 - bdzfl 9-22
- bdzfla 9-22
- bdzflr 9-31
- bdzflrl 9-31
- bdzl 9-21
- bdzla 9-21
- bdzlr 9-31
- bdzlrl 9-31
- bdzt 9-22
- bdzta 9-22
- bdztl 9-22
- bdztla 9-22
- bdztlr 9-31
- bdztlrl 9-31
- beq 9-22
- beqa 9-22
- beqctr 9-27
- beqctrl 9-27
- beql 9-22
- beqlrl 9-31
- bf 9-22
- bfa 9-22
- bfctr 9-27
- bfctrl 9-27
- bfl 9-22
- bfla 9-22
- bflr 9-32
- bflrl 9-32
- bge 9-23
- bgea 9-23
- bgectr 9-27
- bgectrl 9-27
- bgel 9-23
- bgela 9-23
- bgelr 9-32
- bgelrl 9-32
- bgt 9-23
- bgta 9-23
- bgtctr 9-27
- bgtctrl 9-27
- bgtl 9-23
- bgtla 9-23
- bgtlr 9-32
- bgtlrl 9-32
- ble 9-23
- blea 9-23
- blectr 9-27
- blel 9-23
- blela 9-23
- blelr 9-32
- blelrl 9-32
- blr 9-30
- blrl 9-30
- blt 9-23
- blta 9-23
- bltctr 9-27
- bltctrl 9-27
- bltl 9-23
- bltla 9-23
- bltlr 9-32
- bltlrl 9-32
- bne 9-24

bnea 9-24
bnectrl 9-28
bnel 9-24
bnela 9-24
bnelr 9-32
bnelrl 9-32
bng 9-24
bnga 9-24
bngctr 9-28
bngctrl 9-28
bngl 9-24
bnglr 9-32
bnglrl 9-32
bni 9-24
bnla 9-24
bnlctr 9-28
bnll 9-24
bnlla 9-24
bnllr 9-33
bnllrl 9-33
bns 9-24
bnsa 9-24
bnsctr 9-28
bnsctrl 9-28
bnsi 9-24
bnsia 9-24
bnsir 9-33
bnsirl 9-33
bnu 9-25
bnuar 9-25
bnuctr 9-28
bnuctrl 9-28
bnul 9-25
bnula 9-25
bnulr 9-33
bnulrl 9-33
bsalr 9-33
bso 9-25
bsoa 9-25
bsoctr 9-28
bsoctrl 9-28
bsol 9-25
bsola 9-25
bsolrl 9-33
bt 9-25
bta 9-25
btctr 9-28
btctrl 9-28
bti 9-25
btia 9-25
btir 9-33
btirl 9-33
bun 9-25
buna 9-25
bunctr 9-29
bunctrl 9-29
bunl 9-25
bunla 9-25
bunlr 9-33
bunlrl 9-33
crlslwi 9-147
crlwi 9-147
crlwi. 9-147
clrrwi 9-148
clrrwi. 9-148
cmplw 9-36
cmplwi 9-37
cmpw 9-34
cmpwi 9-35
crclr 9-46
crmove 9-44
crnot 9-43
crset 9-41
explained 2-42
extlwi 9-148
extlwi. 9-148
extrwi 9-148
extrwi. 9-148
for addi 9-9
for addic 9-10
for addic. 9-11, 9-115
for addis 9-12
for bc, bca, bcl, bcla 9-21
for bcctr, bcctrl 9-27
for bclr, bclrl 9-30
for cmp 9-34
for cmpi 9-35
for cmpl 9-36
for cmpli 9-37
for creqv 9-41
for crnor 9-43
for cror 9-44
for crxor 9-46
for mfspr 9-113
for mtrcf 9-116
for mtspr 9-120
for nor, nor. 9-139
for or, or. 9-140
for ori 9-142
for rlwimi, rlwimi. 9-146
for rlwinm, rlwinm. 9-147
for rlwnm, rlwnm. 9-150
for subf, subf., subfo, subfo. 9-176
for subfc, subfc., subfco, subfco. 9-177
for tlbre 9-185
for tw 9-191
for twi 9-194
inslwi 9-146
inslwi. 9-146
insrwi 9-146
insrwi. 9-146
li 9-9
lis 9-12
mftb 9-115
mftbu 9-115
mr 9-140
mr. 9-140
mtrc 9-116
nop 9-142
not 9-139
not. 9-139

- rotlw 9-150
- rotlw. 9-150
- rotlwi 9-148
- rotlwi. 9-148
- rotrwi 9-148
- rotrwi. 9-148
- slwi 9-148
- slwi. 9-148
- srwi 9-149
- srwi. 9-149
- sub 9-176
- sub. 9-176
- subc 9-177
- subc. 9-177
- subco 9-177
- subco. 9-177
- subi 9-9
- subic 9-10
- subic. 9-11
- subis 9-12
- subo 9-176
- subo. 9-176
- tblrehi 9-185
- tblrelo 9-185
- tblwehi 9-189
- tblwelo 9-189
- trap 9-191
- tweq 9-191
- tweqi 9-194
- twge 9-191
- twgei 9-194
- twgle 9-191
- twgt 9-191
- twgti 9-194
- twle 9-191
- twlei 9-194
- twlgei 9-194
- twlgt 9-191
- twlgti 9-194
- twlle 9-192
- twllei 9-194
- twllt 9-192
- twllti 9-194
- twlng 9-192
- twlngi 9-194
- twlnl 9-192
- twlnli 9-195
- twlt 9-192
- twlti 9-195
- twne 9-192
- twnei 9-195
- twng 9-192
- twngi 9-195
- twnl 9-192
- twnli 9-195
- extended mnemonics for
 - tlbre 9-189
- external interrupts
 - programming note 5-18
 - register settings 5-19
- extlwi 9-148
- extlwi. 9-148
- extrwi 9-148
- extrwi. 9-148
- extsb 9-63
- extsb. 9-63
- F**
- features
 - DCU 4-1
 - ICU 4-1
- fetcher, improved performance of 4-4
- FIT 6-5
- FIT (fixed interval timer)
 - interrupts, causes 5-23
 - interrupts, register settings 5-24
- fixed interval timer 6-5
- fixed interval timer. *See* FIT
- Fixed Point Exception Register. *See* XER
- G**
- general interrupt handling registers, listed 5-7
- general purpose registers. *See* GPRs
- GPR0-GPR31 10-26
- GPRs (general purpose registers)
 - interrupt control instructions 2-41
 - overview 2-5
 - summary 1-9
- guarded (G) storage attribute
 - controlled by SGR 7-20
 - preventing speculative accesses 2-27
 - virtual mode 7-6
- I**
- I storage attribute. *See* caching inhibited (I) storage attribute
- IAC1-IAC4 10-22, 10-27
- IAC1-IAC4 8-9
- icbi 9-65
 - data storage interrupts 7-17
 - function 4-9
- icbt 9-66
 - data storage interrupts 7-17
 - function 4-9
- iccci 9-67
 - function 4-9
 - when use not recommended 7-17
- ICCR 10-28
- ICCR (Instruction Cache Cachability Register)
 - controlling the I storage attribute 7-20
 - controls cachability 4-5
- ICDBDR 10-30
- ICDBDR (Instruction Cache Debug Data Register)
 - illustrated 4-14, 10-30
 - programming note 4-15
- icread 9-68
 - controlled by CCR0 4-11
 - function 4-9
 - programming note 4-15
- ICU (instruction cache unit)
 - cachability control 4-5
 - cache line fills 4-4
 - coherency 4-6
 - features 4-1

- instruction flow, illustrated 4-4
- instructions 4-9
 - synchronization 4-6
 - synonyms 4-5
- imprecise interrupts 5-1
- initialization
 - code example 3-5
 - of processor 3-3
 - requirements 3-4
 - sequence 3-4
- inslwi 9-146
- inslwi. 9-146
- insrwi 9-146
- insrwi. 9-146
- instruction
 - add 9-6
 - add. 9-6
 - addc 9-7
 - addc. 9-7
 - addco 9-7
 - addco. 9-7
 - adde 9-8
 - adde. 9-8
 - addeo 9-8
 - addeo. 9-8
 - addi 9-9
 - addic 9-10
 - addic. 9-11
 - addis 9-12
 - addme 9-13
 - addme. 9-13
 - addmeo 9-13
 - addmeo. 9-13
 - addo 9-6
 - addo. 9-6
 - addze 9-14
 - addze. 9-14
 - addzeo 9-14
 - addzeo. 9-14
 - and 9-15
 - and. 9-15
 - andc 9-16
 - andc. 9-16
 - andi. 9-17
 - andis. 9-18
 - b 9-19
 - ba 9-19
 - bc 9-20
 - bca 9-20
 - bcctr 9-26
 - bcctrl 9-26
 - bcl 9-20
 - bcla 9-20
 - bclr 9-30
 - bclrl 9-30
 - bl 9-19
 - bla 9-19
 - cmp 9-34
 - cmpi 9-35
 - cmpl 9-36
 - cmpli 9-37
 - cntlzw 9-38
 - cntlzw. 9-38
 - crand 9-39
 - crandc 9-40
 - creqv 9-41
 - crnand 9-42
 - crnor 9-43
 - cror 9-44
 - crorc 9-45
 - crxor 9-46
 - dcbf 9-49
 - dcbi 9-50
 - dcbst 9-51
 - dcbt 9-52
 - dcbz 9-54
 - dccci 9-56
 - dcread 9-57
 - divw 9-59
 - divw. 9-59
 - divwo 9-59
 - divwo. 9-59
 - divwu 9-60
 - divwu. 9-60
 - divwuo 9-60
 - divwuo. 9-60
 - eieio 9-61
 - eqv 9-62
 - eqv. 9-62
 - extsb 9-63
 - extsb. 9-63
 - icbi 9-65
 - icbt 9-66
 - iccci 9-67
 - icread 9-68
 - isync 9-70
 - lbz 9-71
 - lbzu 9-72
 - lbzx 9-74
 - lha 9-75
 - lhau 9-76
 - lhax 9-78
 - lhrbx 9-79
 - lhz 9-80
 - lhzu 9-81
 - lhzux 9-82
 - lhzx 9-83
 - lmw 9-84
 - lswi 9-85
 - lswx 9-87
 - lwarx 9-89
 - lwz 9-91
 - lwzu 9-92
 - lwzux 9-93
 - lwzx 9-94
 - macchw 9-95
 - macchws 9-96
 - macchwsu 9-97
 - macchwu 9-98
 - machhw 9-99
 - machhwsu 9-101
 - machhwu 9-102

maclhw 9-103
 maclhws 9-104, 9-138
 maclhwu 9-106
 mcrf 9-107
 mcrxr 9-108
 mfcrr 9-109
 mfdcr 9-110
 mfmsr 9-111
 mfspr 9-112
 mtcrf 9-116
 mtdcr 9-117
 mtspr 9-119
 mulchw 9-121
 mulchwu 9-122
 mulhhw 9-123
 mulhhwu 9-124
 mulhwu 9-126
 mulhwu. 9-126
 mullhw 9-127
 mullhwu 9-128
 mulli 9-129
 mullw 9-130
 mullw. 9-130
 mullwo 9-130
 mullwo. 9-130
 nand 9-131
 nand. 9-131
 neg 9-132
 neg. 9-132
 nego 9-132
 nego. 9-132
 nmacchw 9-133
 nmacchws 9-134
 nmachhw 9-135
 nmachhws 9-136
 nmaclhw 9-137
 nmaclhws 9-138
 nor 9-139
 nor. 9-139
 or 9-140
 or. 9-140
 orc 9-141
 orc. 9-141
 ori 9-142
 oris 9-143
 rfcir 9-144
 rfi 9-145
 rlwimi 9-146
 rlwimi. 9-146
 rlwinm 9-147
 rlwinm. 9-147
 rlwnm 9-150
 rlwnm. 9-150
 sc 9-151
 slw 9-152
 slw. 9-152
 saw 9-153
 saw. 9-153
 sawi 9-154
 sawi. 9-154
 saw 9-155
 srw. 9-155
 stb 9-156
 stbu 9-157
 stbux 9-158
 stbx 9-159
 sth 9-160
 sthbrx 9-161
 sthu 9-162
 sthux 9-163
 sthx 9-164
 stmw 9-165
 stswi 9-166
 stswx 9-167
 stw 9-169
 stwbrx 9-170
 stwcx. 9-171
 stwu 9-173
 stwux 9-174
 stwx 9-175
 subf 9-176
 subf. 9-176
 subfc 9-177
 subfc. 9-177
 subfco 9-177
 subfco. 9-177
 subfe 9-178
 subfe. 9-178
 subfeo 9-178
 subfeo. 9-178
 subfic 9-179
 subfme 9-180
 subfme. 9-180
 subfmeo 9-180
 subfmeo. 9-180
 subfo 9-176
 subfo. 9-176
 subfze 9-181
 subfze. 9-181
 subfzeo 9-181
 subfzeo. 9-181
 sync 9-182
 tlbia 9-183
 tlbre 9-184
 tlbsx 9-186
 tlbsx. 9-186
 tlbsync 9-187
 tlbwe 9-188
 tw 9-190
 twi 9-193
 wrtee 9-196
 wrteei 9-197
 xor 9-198
 xori 9-199
 instruction cache array, improved fetcher performance 4-4
 Instruction Cache Cachability Register. *See* ICCR
 Instruction Cache Cachability Register. *See* ICCR
 Instruction Cache Debug Data Register. *See* ICDBDR
 instruction cache synonym, defined 4-5
 instruction cache unit. *See* ICU
 instruction fetching

- from little endian storage 2-20
- instruction fields A-41
- instruction formats 9-2, A-41
 - diagrams A-43
- instruction forms A-41, A-43
- instruction queue
 - illustrated 2-24
 - role in processing instructions 2-23
- instruction set
 - brief summaries by category 2-36
 - for embedded processors 2-37
- instruction set portability 9-1
- instruction set summary
 - cache control 2-41
 - CR logical 2-40
- instruction storage interrupts
 - causes 5-17
 - register settings 5-18
- instruction timings C-3
 - branches and cr logicals C-3
 - general rules C-3
 - instruction cache misses C-7
 - loads and stores C-7
 - strings C-6
- instruction TLB. *See* ITLB
- instructions
 - alphabetical, including extended mnemonics A-1
 - arithmetic and logical 2-38, B-33
 - arithmetic compares 2-11
 - branch 2-40, B-38
 - branch conditional, testing CTR 2-25
 - byte-reverse, usefulness of 2-21
 - cache
 - DAC debug events 8-15
 - cache control B-41
 - cache control, alignment of 2-16
 - compare 2-39
 - comparison B-39
 - condition register logical B-37
 - context synchronizing, defined 2-33
 - CR logical 2-39
 - extended mnemonics B-9
 - format diagrams A-43
 - formats A-41
 - forms A-41, A-43
 - ICU controlling 4-9
 - interrupt control 2-41, B-42
 - logical compares 2-11
 - opcodes A-33
 - privileged B-7
 - privileged, listed 2-31
 - processor management 2-42, B-44
 - for reading DCRs 2-32
 - for reading privileged SPRs 2-32
 - rotate and shift B-40
 - specific to PowerPC Embedded Controllers B-5
 - storage reference B-29
 - storage reference, alignment of 2-16
 - storage reference, in core 2-37
 - TLB management 2-42, B-42
- interfaces
 - interrupt controller 1-8
 - interrupt controller interface 1-8
- interrupts
 - alignment 2-17
 - register settings 5-19
 - summary 5-19
 - asynchronous, defined 5-1
 - behavior 5-1
 - critical
 - defined 5-5
 - processing 5-6
 - critical input 5-13
 - data machine check 5-15
 - data storage 5-16, 7-10
 - register settings 5-17
 - debug, register settings 5-26
 - defined 5-1
 - DTLB miss 7-11
 - DTLB, register settings 5-25
 - external
 - programming note 5-18
 - register settings 5-19
 - fetching past, speculatively 2-28
 - FIT, causes 5-23
 - FIT, register settings 5-24
 - handling as critical 5-3
 - handling priorities 5-3
 - handling priorities, illustrated 5-4
 - imprecise, defined 5-1
 - instruction storage 7-10
 - causes 5-17
 - register settings 5-18
 - ITLB miss 7-11
 - ITLB miss, registers 5-25
 - machine check, causes of 5-14
 - machine check, defined 5-2
 - machine check—instruction
 - handling 5-14
 - register settings 5-15
 - synchronism 5-3
 - noncritical
 - defined 5-5
 - processing 5-5
 - PIT, register settings 5-22
 - precise handling 5-2
 - precise, defined 5-1
 - program 7-11
 - causes 5-20
 - ESR usage 5-20
 - register settings 5-21
 - register settings during critical 5-14
 - synchronous, defined 5-2
 - system call, register settings 5-22
 - TLB miss, preventing 7-11
 - TLB-related 7-9
 - vector offsets, illustrated 5-6
 - WDT, causes 5-24
 - WDT, register settings 5-24
- isync 9-70
 - and ITLB 7-7
 - context synchronization, example 2-35

ITLB (instruction translation lookaside buffer)
accesses 7-7
consistency 7-7
defined 7-6
miss interrupts 5-25, 7-11
programming note 7-8

L

lbz 9-71
lbzu 9-72
lbzx 9-74
lha 9-75
lhau 9-76
lhax 9-78
lhbrx 9-79
lhz 9-80
lhzu 9-81
lhzux 9-82
lhzx 9-83
li 9-9
Link Register. *See* LR
lis 9-12
little endian
alignment 2-17
byte ordering supported 2-19
defined 2-18
mapping 2-19
storage attributes 2-20
storage regions
accessing data from 2-21
byte-reverse instructions 2-21- 2-23
fetching instructions from 2-20
lmw 9-84
load strategies, controlled by DCU 4-8
logical compares 2-11
logical instructions
CR 2-39
overview 2-38
LR 10-31
LR (Link Register)
branch instructions 2-40
function 2-7
lswi 9-85
lswx 9-87
lwarx 9-89
lwz 9-91
lwzu 9-92
lwzux 9-93
lwzx 9-94

M

macchw 9-95
macchws 9-96
macchwsu 9-97
macchwu 9-98
machhw 9-99
machhwsu 9-101
machhwu 9-102
machine check interrupts
causes 5-14
defined 5-2
machine check—instruction interrupts

handling 5-14
register settings 5-15
synchronism 5-3
Machine State Register. *See* MSR
maclhw 9-103
maclhws 9-104, 9-138
maclhwu 9-106
mapping
big endian 2-19
little endian 2-19
structure, examples 2-18
mcrf 9-107
mcrxr 9-108
memory mapping
of hardware 2-29
memory models, non-supported 4-8, 7-5
memory organization 2-1
mfcr 9-109
mfocr 9-110
mfmsr 9-111
mfspr 9-112
mftb 9-115
mftbu 9-115
misalignments, defined 2-17
MMU (memory management unit)
accesses, interrupts from 7-10, 7-11
address translation 7-1
data storage interrupts 7-10
DTLB miss interrupts 7-11
execute permissions 7-14
general access protection 7-13
instruction storage interrupts 7-10
ITLB miss interrupts 7-11
MSR and access protection 7-13
overview 1-5
program interrupts 7-11
recording page references and changes 7-12
TLB management 7-11
zone protection 7-14
mnemonics, extended. *See* extended mnemonics
modes
execution 2-31
real, storage attribute control 7-17
mr 9-140
mr. 9-140
MSR 2-13, 10-32
MSR (Machine State Register)
bits and exception handling 2-31
contents after resets 3-2
controlling execution mode 2-31
DR bit 7-1
illustrated 5-7
interrupt control instructions 2-41
IR bit 7-1
programming note 5-7
summarized 1-9
mtcr 9-116
mtcrf 9-116
mtocr 9-117
mtmsr
execution synchronization 2-35

mtspr 9-119
mulchw 9-121
mulchwu 9-122
mulhhw 9-123
mulhhwu 9-124
mulhwu 9-126
mulhwu. 9-126
mullhw 9-127
mullhwu 9-128
mulli 9-129
mullw 9-130
mullw. 9-130
mullwo 9-130
mullwo. 9-130

N

nand 9-131
nand. 9-131
neg 9-132
neg. 9-132
nego 9-132
nego. 9-132
nmacchw 9-133
nmacchws 9-134
nmachhw 9-135
nmachhws 9-136
nmaclhw 9-137
nmaclhws 9-138
noncritical interrupts
 defined 5-5
 processing 5-5
nop 9-142
nor 9-139
nor. 9-139
not 9-139
not. 9-139
notation xxii, 9-2, A-41
notational conventions xxii

O

opcodes A-33
optimization
 coding guidelines C-1
 alignment C-2
 boolean variables C-1
 branch prediction C-2
 dependency upon CR C-2
 floating point emulation C-1
or 9-140
or. 9-140
orc 9-141
orc. 9-141
ori 9-142
oris 9-143

P

page identification fields, UTLB 7-3
performance
 DCU
 improve with simultaneous caching 4-17
 limited by sequential caching 4-18
 overview 4-16
 improve

 through byte-writeability 4-6
 lower
 from cache-inhibited regions 4-5
PID 10-34
PID (process ID)
 illustrated 7-13
PIT 6-4, 10-35
PIT (programmable interval timer)
 interrupts, register settings 5-22
portability, instruction set 9-1
PowerPC architecture 1-3
precise interrupts 5-1
pre-fetch
 branches to CTR 2-28
 branches to LR 2-28
 buffers 2-23
 past interrupts 2-28
primary opcodes A-33
priority signal
 DCU 4-17
privileged mode
 defined 2-30
 instructions, listed 2-31
 registers 2-4
privileged programming model 2-1
privileged SPRs
 instructions for reading 2-32
problem state. *See* user mode
process ID. *See* PID
processor
 management instructions 2-42
Processor Version Register. *See* PVR
program interrupts
 and TLB 7-11
 causes 5-20
 ESR usage 5-20
 programming note 7-11
 register settings 5-21
programmable interval timer 6-4
programming model
 features 2-1
programming models
 privileged 2-1
 user 2-1
programming note
 data storage interrupts 5-16
 EA access in DCU 4-9
 external or timer interrupts 5-18
 instruction pipeline 4-15
 MSR affected by instructions 5-7
 non-supported memory models 4-8
 program interrupts 7-11
 reserved fields 2-2
 RPN field 7-4
 synchronizing the ITLB 7-8
pseudocode 9-2
PVR 10-36
PVR (Processor Version Register)
 illustrated 2-10

R

real mode

- storage attribute control 7-17
- register set summary 1-9
- registers
 - categories 2-2, 10-1
 - CCR0 10-6
 - CR 10-1, 10-8
 - CTR 10-9
 - DAC1 8-9
 - DAC1–DAC2 8-9, 10-10
 - DBCR 8-4
 - DBCR0 10-11, 10-13
 - DBSR 8-7, 10-15
 - DCCR 10-17
 - DCR numbering 10-4
 - DCRs
 - summarized 1-9
 - uses for 2-15
 - DCWR 10-19
 - DEAR 10-21
 - descriptions of commonly used 2-2
 - during debug exceptions 5-26
 - ESR 10-23
 - EVPR 10-25
 - GPR 10-1
 - GPR0–GPR31 10-26
 - GPRs
 - overview 2-5
 - summary 1-9
 - IAC1–IAC4 10-22, 10-27
 - IAC1–IAC4 8-9
 - ICCR 10-28
 - ICDBDR 10-30
 - interrupt handling 5-7
 - LR 10-31
 - MSR 2-13, 10-1, 10-32
 - PID 10-34
 - PIT 6-4, 10-35
 - PVR 10-36
 - reserved 10-1
 - reserved fields 2-2, 10-1
 - SGR 10-37
 - SLER 10-39
 - SPR numbering 10-2
 - SPRG0–SPRG4 2-10
 - SPRG0–SPRG7 10-41
 - SPRs
 - overview 2-5
 - summary 1-9
 - SRR0 10-42
 - SRR1 10-43
 - SRR2 10-44
 - SRR3 10-45
 - SU0R 10-46
 - summary of sets 1-9
 - supervisor, illustrated 2-4
 - TBL 10-48
 - TBU 10-49
 - TCR 6-5, 6-6, 6-9, 10-50
 - TSR 6-6, 6-8, 10-51
 - user, illustrated 2-4
 - USPRG0 2-10, 10-52
 - XER 10-53
 - ZPR 10-54
- reservation bit 9-89, 9-171
- reserved fields 10-1
 - programming note 2-2
- reserved registers 10-1
- resets
 - effects on MSR 3-2
 - effects on SPRs 3-3
 - processor initialization 3-3
 - processor state after 3-1
- rftci 9-144
 - effect on MSR reserved fields 5-7
- rfti 9-145
 - effect on MSR reserved fields 5-7
- rlwimi 9-146
- rlwimi. 9-146
- rlwinm 9-147
- rlwinm. 9-147
- rlwnm 9-150
- rlwnm. 9-150
- rotlw 9-150
- rotlw. 9-150
- rotlwi 9-148
- rotlwi. 9-148
- rotrwi 9-148
- rotrwi. 9-148
- extended mnemonics
 - bnectr 9-28

S

- Save/Restore Registers 0-1. *See* SRR0-1
- sc 9-151
- secondary opcodes A-33
- SGR 10-37
- SGR (Storage Guarded Register)
 - controlling speculative accesses 2-27
 - controlling the guarded (G) storage attribute 7-20
- shadow TLB. *See* DTLB
- SLER 10-39
- SLER (Storage Little Endian Register)
 - controlling the endian (E) storage attribute 7-20
- slw 9-152
- slw. 9-152
- slwi 9-148
- slwi. 9-148
- Special Purpose Register General 0-7. *See* SPRG0-7
- special purpose registers. *See* SPRs
- speculative accesses
 - to CTR or LR 2-28
 - defined 2-27
 - down predicted path 2-28
 - fetching past interrupts 2-28
 - fetching past tw or twi 2-29
 - fetching past unconditional branches 2-29
 - preventing inappropriate 2-27- 2-30
- SPRG0-7 (Special Purpose Register General 0-7)
 - temporary storage to 2-9
- SPRG0–SPRG4 2-10
- SPRG0–SPRG7 10-41
- SPRs (special purpose registers)
 - contents after resets 3-3

- listed, with page references 2-6
- overview 2-5
- privileged and non-privileged 2-5
- privileged, instructions for reading 2-32
- summary 1-9
- sraw 9-153
- sraw. 9-153
- srawi 9-154
- srawi. 9-154
- SRR0 10-42
- SRR0-1 (Save/Restore Registers 0-1)
 - illustrated 5-9
- SRR1 10-43
- SRR2 10-44
- SRR3 10-45
- srw 9-155
- srw. 9-155
- srwi 9-149
- srwi. 9-149
- stb 9-156
- stbu 9-157
- stbux 9-158
- stbx 9-159
- sth 9-160
- sthbrx 9-161
- sthu 9-162
- sthux 9-163
- sthx 9-164
- stmw 9-165
- storage attribute control registers
 - DCCR 7-20
 - DCWR 7-19
 - ICCR 7-20
 - SGR 7-20
 - SLER 7-20
 - SUOR 7-20
- storage attributes
 - caching inhibited (I)
 - real mode 7-20
 - virtual mode 7-5
 - endian (E)
 - and little endian 2-19
 - real mode 7-20
 - when controlled by TLB 7-6
 - guarded (G)
 - controlling speculative accesses 2-27
 - real mode 7-20
 - virtual mode 7-6
 - memory coherent (M)
 - not supported 7-6
 - overview 2-2
 - real mode 7-17
 - TLB control of 7-5
 - user-defined (U0)
 - real mode 7-20
 - virtual mode 7-6
 - virtual mode 7-5
 - write-through (W)
 - real mode 7-19
 - virtual mode 7-5
- Storage Guarded Register. *See* SGR
- Storage Little Endian Register. *See* SLER
- storage reference instructions 2-37
- storage regions
 - big endian
 - alignment 2-17
 - byte-reverse instructions 2-21- 2-22
 - little endian
 - accessing data from 2-21
 - alignment 2-17
 - byte reversal 2-21
 - byte-reverse instructions 2-21- 2-23
 - data alignment 2-21
 - fetching instructions from 2-20
- storage synchronization 2-35
- Storage User-Defined 0 Register. *See* SUOR
- string instructions
 - access protection 7-17
- structure mapping
 - examples 2-18
- stswi 9-166
- stswx 9-167
- stw 9-169
- stwbrx 9-170
- stwcx. 9-171
- stwu 9-173
- stwux 9-174
- stwx 9-175
- SUOR 10-46
- SUOR (Storage User-Defined 0 Register)
 - controlling the user-defined (U0) storage attribute 7-20
- sub 9-176
- sub. 9-176
- subc 9-177
- subc. 9-177
- subco 9-177
- subco. 9-177
- subf 9-176
- subf. 9-176
- subfc 9-177
- subfc. 9-177
- subfco 9-177
- subfco. 9-177
- subfe 9-178
- subfe. 9-178
- subfeo 9-178
- subfeo. 9-178
- subfic 9-179
- subfme 9-180
- subfme. 9-180
- subfmeo 9-180
- subfmeo. 9-180
- subfo 9-176
- subfo. 9-176
- subfze 9-181
- subfze. 9-181
- subfzeo 9-181
- subfzeo. 9-181
- subi 9-9
- subic 9-10

- subic. 9-11
- subis 9-12
- subo 9-176
- subo. 9-176
- supervisor state. *See* privileged mode
- sync 9-182
 - storage synchronization 2-35
- synchronization
 - context 2-33
 - execution, defined 2-35
 - ICU 4-6
 - references to PowerPC Architecture 2-33
 - storage 2-35
- synchronous interrupts 5-2
- system call interrupts
 - register settings 5-22

T

- TBL 10-48
- tblrehi 9-185
- tblrelo 9-185
- tblwehi 9-189
- tblwelo 9-189
- TBU 10-49
- TCR 6-9, 10-50
- TID (translation ID)
 - and MMU access protection
- time base 6-1
 - implementation 2-13
 - writing 2-13
- timer interrupts
 - programming note 5-18
- timers
 - FIT 6-5
 - fixed interval timer 6-5
 - PIT 6-4
 - programmable interval timer 6-4
 - TCR 6-9
 - timer control register 6-9
 - timer status register 6-8
 - TSR 6-8
 - watchdog 6-6
- timings
 - instruction C-3
 - branches and cr logicals C-3
 - general rules C-3
 - instruction cache misses C-7
 - loads and stores C-7
 - strings C-6
- TLB (translation lookaside buffer) 7-2
 - access protection 7-13, 7-16
 - and cacheability control 4-8
 - execute permissions 7-14
 - interrupts 7-9
 - invalidate instruction 7-12
 - management instructions 2-42
 - preventing miss interrupts 7-11
 - read/write instructions 7-12
 - search instructions 7-12
 - sync instruction 7-12
 - zone protection 7-14
 - See also* ITLB;UTLB;DTLB
- tlbia 9-183
 - and TLB management 7-12
- tlbre 9-184
 - and TLB management 7-12
- tlbsx 9-186
 - and TLB management 7-12
- tlbsx. 9-186
 - and TLB management 7-12
- tlbsync 9-187
 - and TLB management 7-12
- tlbwe 9-188
 - and TLB management 7-12
- translation ID. *See* TID
- translation lookaside buffer. *See* TLB
- translation, address. *See* address translation
- trap 9-191
- TSR 6-8, 10-51
- tw 9-190
 - fetching past 2-29
- tweq 9-191
- tweqi 9-194
- twge 9-191
- twgei 9-194
- twgle 9-191
- twgt 9-191
- twgti 9-194
- twi 9-193
 - fetching past 2-29
- twle 9-191
- twlei 9-194
- twlgei 9-194
- twlgt 9-191
- twlgti 9-194
- twlle 9-192
- twllei 9-194
- twllt 9-192
- twllti 9-194
- twlng 9-192
- twlngi 9-194
- twlnl 9-192
- twlnli 9-195
- twlt 9-192
- twlti 9-195
- twne 9-192
- twnei 9-195
- twng 9-192
- twngi 9-195
- twnl 9-192
- twnli 9-195

U

- unconditional branches
 - AA field 2-24
 - speculative accesses 2-29
- unified TLB. *See* UTLB
- user mode
 - defined 2-30
 - registers 2-4
- user programming model 2-1
- user-defined (U0) storage attribute
 - controlled by SU0R 7-20
 - virtual mode 7-6

- USPRG0 2-10, 10-52
- UTLB (unified translation lookaside buffer)
 - access control fields 7-5
 - entry format, illustrated 7-3
 - EPN field 7-3
 - EX field 7-5
 - field categories 7-3
 - functional overview 7-2
 - page identification fields 7-3
 - RPN field 7-4
 - SIZE field 7-4
 - TID field 7-4
 - translation field 7-4
 - V field 7-4
 - WR field 7-5
 - ZSEL field 7-5

V

- virtual mode
 - and TLB control of storage attributes 7-5

W

- watchdog timer 6-6
- WDT (watchdog timer)
 - interrupts, causes 5-24
 - interrupts, register settings 5-24
- write strategies
 - controlled by DCWR 4-7
 - used by DCU 4-7
- write-through (W) storage attribute
 - controlled by DCWR 7-19
 - when controlled by TLB 7-5
- wrtee 9-196
- wrteei 9-197

X

- XER 10-53
- XER (Fixed Point Exception Register)
 - illustrated 2-7
- xor 9-198
- xori 9-199

Z

- zone fault 5-16
- Zone Protection Register. *See* ZPR
- zone, defined 7-14
- ZPR 10-54
- ZPR (Zone Protection Register)
 - illustrated 7-14



© International Business Machines Corporation 1996, 2001
Printed in the United States of America
12/3/01
All Rights Reserved

The information contained in this document is subject to change without notice. The products described in this document are NOT intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change IBM's product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of IBM or third parties. All information contained in this document was obtained in specific environments, and is presented as illustration. The results obtained in other operating environments may vary.

THE INFORMATION CONTAINED IN THIS DOCUMENT IS PROVIDED ON AN "AS IS" BASIS. In no event will IBM be liable for any damages arising directly or indirectly from any use of the information contained in this document.

IBM Microelectronics Division
1580 Route 52, Bldg. 504
Hopewell Junction, NY
12533-6531

The IBM home page can be found at www.ibm.com

The IBM Microelectronics Division home page can be found at www.chips.ibm.com

Address technical queries about this product to ppcsupp@us.ibm.com