ADSP-BF561 Blackfin® Processor Hardware Reference

Revision 1.2, January 2010

Part Number
82-000561-01

Analog Devices, Inc.
One Technology Way
Norwood, Mass. 02062-9106

ANALOG DEVICES

Copyright Information

© 2010 Analog Devices, Inc., ALL RIGHTS RESERVED. This document may not be reproduced in any form without prior, express written consent from Analog Devices, Inc.

Printed in the USA.

Disclaimer

Analog Devices, Inc. reserves the right to change this product without prior notice. Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use; nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under the patent rights of Analog Devices, Inc.

Trademark and Service Mark Notice

The Analog Devices logo, Blackfin, the Blackfin logo, CrossCore, EZ-KIT Lite, SHARC, TigerSHARC, and VisualDSP++ are registered trademarks of Analog Devices, Inc.

All other brand and product names are trademarks or service marks of their respective owners.

CONTENTS

PREFACE

Purpose of This Manual xxxv
Intended Audience xxxv
Manual Contents xxxvi
What's New in This Manual xxxix
Technical or Customer Support xl
Supported Processors xl
Product Information xli
Analog Devices Web Site xli
VisualDSP++ Online Documentation xli
Technical Library CD xlii
Conventions xliii
Register Diagram Conventions xlv

INTRODUCTION

ADSP-BF561 Peripherals 1-1
ADSP-BF561 Core Architecture 1-3

Memory Architecture 1-6
Internal Memory 1-8
External Memory 1-8
I/O Memory Space 1-9
Event Handling 1-9
Core Event Controller (CEC) 1-10
System Interrupt Controller (SIC) 1-11
DMA Support 1-11
External Bus Interface Unit 1-12
PC133 SDRAM Controller 1-12
Asynchronous Controller 1-13
Parallel Peripheral Interface 1-13
Serial Ports (SPORTs) 1-15
Serial Peripheral Interface (SPI) Port 1-16
Timers 1-17
UART Port 1-18
Watchdog Timer 1-19
Programmable Flags 1-20
Clock Signals 1-21
Dynamic Power Management 1-21
Full On Operating Mode (Maximum Performance) 1-22
Active Operating Mode (Moderate Power Savings) 1-22
Sleep Operating Mode (High Power Savings) 1-22

Deep Sleep Operating Mode (Maximum Power Savings) 1-23
Hibernate State 1-23
Voltage Regulation 1-23
Boot Modes 1-24
Instruction Set Description 1-25
Development Tools 1-26

COMPUTATIONAL UNITS

Using Data Formats 2-3
Binary String 2-3
Unsigned 2-4
Signed Numbers: Two's-Complement 2-4
Fractional Representation: 1.15 2-4
Register Files 2-5
Data Register File 2-6
Accumulator Registers 2-6
Pointer Register File 2-6
DAG Register Set 2-7
Register File Instruction Summary 2-8
Data Types 2-10
Endianess 2-12
ALU Data Types 2-12
Multiplier Data Types 2-13
Shifter Data Types 2-14
Arithmetic Formats Summary 2-15

Using Multiplier Integer and Fractional Formats 2-16
Rounding Multiplier Results 2-17
Unbiased Rounding 2-18
Biased Rounding 2-20
Truncation 2-21
Special Rounding Instructions 2-22
Using Computational Status 2-22
Arithmetic Status Register (ASTAT) 2-23
Arithmetic Logic Unit (ALU) 2-23
ALU Operations 2-23
Single 16-Bit Operations 2-25
Dual 16-Bit Operations 2-25
Quad 16-Bit Operations 2-26
Single 32-Bit Operations 2-27
Dual 32-Bit Operations 2-28
ALU Instruction Summary 2-28
ALU Data Flow Details 2-33
Dual 16-Bit Cross Options 2-34
ALU Status Signals 2-35
ALU Division Support Features 2-36
Special SIMD Video ALU Operations 2-36

Multiply Accumulators (Multipliers) 2-36
Multiplier Operation 2-37
Placing Multiplier Results in Multiplier Accumulator Registers 2-38
Rounding or Saturating Multiplier Results 2-38
Saturating Multiplier Results on Overflow 2-39
Multiplier Instruction Summary 2-39
Multiplier Instruction Options 2-41
Multiplier Data Flow Details 2-42
Multiply Without Accumulate 2-44
Special 32-Bit Integer MAC Instruction 2-46
Dual MAC Operations 2-47
Barrel Shifter (Shifter) 2-48
Shifter Operations 2-48
Two-Operand Shifts 2-49
Immediate Shifts 2-49
Register Shifts 2-50
Three-Operand Shifts 2-50
Immediate Shifts 2-50
Register Shifts 2-51
Bit Test, Set, Clear, Toggle 2-52
Field Extract and Field Deposit 2-52
Shifter Instruction Summary 2-52

OPERATING MODES AND STATES

User Mode 3-3
Protected Resources and Instructions 3-4
Protected Memory 3-5
Entering User Mode 3-5
Example Code to Enter User Mode Upon Reset 3-5
Return Instructions That Invoke User Mode 3-5
Supervisor Mode 3-6
Non-OS Environments 3-7
Example Code to Stay in Supervisor Mode Coming Out of Reset 3-8
Emulation Mode 3-9
Idle State 3-9
Example Code for Transition to Idle State 3-10
Reset State 3-10
System Reset and Powerup Configuration 3-12
Hardware Reset 3-13
System Reset Configuration Register A (SICA_SYSCR) 3-14
System Reset Configuration Register B (SICB_SYSCR) 3-15
Software Resets and Watchdog Timer 3-16
Software Reset Register (SICA_SWRST) 3-17
Core-Only Software Reset 3-18
Core and System Reset 3-18

Booting Methods 3-19
SPI Slave Mode Boot Through Master Host (BMODE = 10) 3-21

PROGRAM SEQUENCER

Sequencer-Related Registers 4-3
Sequencer Status Register (SEQSTAT) 4-4
Zero-Overhead Loop Registers (LC, LT, LB) 4-4
System Configuration Register (SYSCFG) 4-6
Instruction Pipeline 4-6
Branches and Sequencing 4-9
Direct Short and Long Jumps 4-10
Direct Call 4-11
Indirect Branch and Call 4-11
PC-Relative Indirect Branch and Call 4-12
Condition Code Flag 4-12
Conditional Branches 4-13
Conditional Register Move 4-13
Branch Prediction 4-14
Loops and Sequencing 4-15
Events and Sequencing 4-18
System Interrupt Processing 4-20
System Peripheral Interrupts 4-22
System Interrupt Wakeup-Enable Registers (SIC_IWRx) 4-26
System Interrupt Status Registers (SIC_ISRx) 4-27

System Interrupt Mask Registers (SIC_IMASKx) 4-29
System Interrupt Assignment Registers (SIC_IARx) 4-30
Core Event Controller Registers 4-36
Core Interrupt Mask Register (IMASK) 4-36
Core Interrupt Latch Register (ILAT) 4-37
Core Interrupt Pending Register (IPEND) 4-38
Global Enabling/Disabling of Interrupts 4-39
Event Vector Table 4-40
Emulation 4-41
Reset 4-41
NMI (Non-Maskable Interrupt) 4-43
Exceptions 4-43
Exceptions While Executing an Exception Handler 4-48
Hardware Error Interrupt 4-49
Core Timer 4-50
General-Purpose Interrupts (IVG7-IVG15) 4-50
Servicing Interrupts 4-51
Nesting of Interrupts 4-52
Non-Nested Interrupts 4-52
Nested Interrupts 4-53
Example Prolog Code for Nested Interrupt Service Routine 4-55
Example Epilog Code for Nested Interrupt Service Routine 4-55
Logging of Nested Interrupt Requests 4-56

Exception Handling 4-57
Deferring Exception Processing 4-57
Example Code for an Exception Handler 4-58
Example Code for an Exception Routine 4-60
Example Code for Using Hardware Loops in an ISR 4-60
Additional Usability Issues 4-61
Executing RTX, RTN, or RTE in a Lower Priority Event 4-61
Recommendation for Allocating the System Stack 4-62
Latency in Servicing Events 4-62

DATA ADDRESS GENERATORS

Addressing With DAGs 5-4
Frame and Stack Pointers 5-5
Addressing Circular Buffers 5-6
Addressing With Bit-reversed Addresses 5-9
Indexed Addressing With Index and Pointer Registers 5-9
Auto-Increment and Auto-Decrement Addressing 5-10
Pre-Modify Stack Pointer Addressing 5-11
Indexed Addressing With Immediate Offset 5-11
Post-Modify Addressing 5-11
Modifying DAG and Pointer Registers 5-12
Memory Address Alignment 5-13
DAG Instruction Summary 5-16

MEMORY

Memory Architecture 6-1
Internal Memory 6-4
Overview of L1 Instruction SRAM 6-5
Overview of L1 Data SRAM 6-5
Overview of Scratchpad Data SRAM 6-6
Overview of On-Chip L2 Memory 6-6
L1 Instruction Memory 6-7
Instruction Memory Control Register (IMEM_CONTROL) 6-7
L1 Instruction SRAM 6-9
L1 Instruction Cache 6-11
Cache Lines 6-12
Cache Hits and Misses 6-14
Cache-Line Fills 6-15
Line-Fill Buffer 6-16
Cache-Line Replacement 6-16
Instruction Cache Management 6-17
Instruction Cache Locking by Line 6-17
Instruction Cache Locking by Way 6-18
Instruction Cache Invalidation 6-19

Instruction Test Registers 6-20
ITEST_COMMAND Register 6-21
ITEST_DATA1 Register 6-22
ITEST_DATA0 Register 6-23
L1 Data Memory 6-24
Data Memory Control Register (DMEM_CONTROL) 6-24
L1 Data SRAM 6-26
L1 Data Cache 6-30
Example of Mapping Cacheable Address Space into Data Banks 6-31
Data Cache Access 6-34
Cache Write Method 6-36
Write Buffers 6-36
Interrupt Priority Register (IPRIO) and Write Buffer Depth 6-37
Data Cache Control Instructions 6-38
Data Cache Invalidation 6-39
Data Test Registers 6-39
Data Test Command Register (DTEST_COMMAND) 6-41
Data Test Data 1 Register (DTEST_DATA1) 6-42
Data Test Data 0 Register (DTEST_DATA0) 6-42
On-Chip Level 2 (L2) Memory 6-44
On-Chip L2 Bank Access 6-44
Latency 6-44
External Memory 6-47

Memory Protection and Properties 6-47
Memory Management Unit 6-47
Memory Pages 6-49
Memory Page Attributes 6-49
Page Descriptor Table 6-50
CPLB Management 6-51
MMU Application 6-53
Examples of Protected Memory Regions 6-55
ICPLB Data Registers (ICPLB_DATAx) 6-56
DCPLB Data Registers (DCPLB_DATAx) 6-57
DCPLB Address Registers (DCPLB_ADDRx) 6-60
ICPLB Address Registers (ICPLB_ADDRx) 6-61
DCPLB and ICPLB Status Registers (DCPLB_STATUS, ICPLB_STATUS) 6-62
DCPLB and ICPLB Fault Address Registers (DCPLB_FAULT_ADDR, ICPLB_FAULT_ADDR) 6-64
Memory Transaction Model 6-66
Load/Store Operation 6-67
Interlocked Pipeline 6-67
Ordering of Loads and Stores 6-68
Synchronizing Instructions 6-69
Speculative Load Execution 6-70
Conditional Load Behavior 6-71

Working With Memory 6-72
Alignment 6-72
Cache Coherency 6-72
Atomic Operations 6-73
Memory-Mapped Registers 6-73
Core MMR Programming Code Example 6-74
Terminology 6-75

CHIP BUS HIERARCHY

Internal Interfaces 7-1
Internal Clocks 7-5
Core Overview 7-5
Core Interfaces 7-5
System Overview 7-7
P Port Interface 7-7
D Port Interface 7-8
On-Chip L2 Memory Interface 7-10
Peripheral Access Bus (PAB) 7-14
PAB Arbitration 7-14
PAB Performance 7-15
DMA-Related Buses 7-16
Peripheral DMA 7-18
DAB, DCB, and DEB Performance 7-18
External Access Bus (EAB) 7-20
EAB Bus and DMA Bus Access Arbitration to EBIU 7-21

DYNAMIC POWER MANAGEMENT

Clocking 8-1
Phase-Locked Loop and Clock Control 8-2
PLL Overview 8-2
PLL Clock-Multiplier Ratios 8-3
Core Clock/System Clock Ratio Control 8-4
PLL Memory-Mapped Registers (MMRs) 8-6
PLL Divide Register (PLL_DIV) 8-6
PLL Control Register (PLL_CTL) 8-7
PLL Status Register (PLL_STAT) 8-9
PLL Lock Count Register (PLL_LOCKCNT) 8-11
Dynamic Power Management Controller 8-11
Operating Modes 8-12
Dynamic Power Management Controller States 8-13
Full On Mode 8-13
Active Mode 8-14
Sleep Mode 8-14
Deep Sleep Mode 8-14
Hibernate State 8-15
Operating Mode Transitions 8-15
Programming Operating Mode Transitions 8-18
PLL Programming Sequence 8-19
PLL Programming Sequence Continues 8-21
Examples 8-21

Active Mode to Full On Mode 8-21
Full On Mode to Active Mode 8-22
In the Full On Mode, Change CLKIN to VCO Multiplier From 31x to 2x 8-22
Dynamic Supply Voltage Control 8-23
Power Supply Management 8-23
Voltage Regulator Control Register (VR_CTL) 8-24
Changing Voltage 8-27
Powering Down the Core (Hibernate State) 8-28

DIRECT MEMORY ACCESS

DMA and Memory DMA MMRs 9-4
Naming Conventions for DMA MMRs 9-6
Naming Conventions for Memory DMA Registers 9-8
Next Descriptor Pointer Register (DMAx_y_NEXT_DESC_PTR/ MDMA_yy_NEXT_DESC_PTR) 9-10
Start Address Register (DMAx_y_START_ADDR/ MDMAx_yy_START_ADDR) 9-11
DMA Configuration Register (DMAx_y_CONFIG/MDMAx_yy_CONFIG) 9-12
Inner Loop Count Register (DMAx_y_X_COUNT/MDMAx_yy_X_COUNT) 9-16
Inner Loop Address Increment Register (DMAx_y_X_MODIFY/MDMAx_yy_X_MODIFY) 9-16

Outer Loop Count Register (DMAx_y_Y_COUNT/MDMAx_yy_Y_COUNT) 9-17
Outer Loop Address Increment Register (DMAx_y_Y_MODIFY/MDMAx_yy_Y_MODIFY) 9-18
Current Descriptor Pointer Register (DMAx_y_CURR_DESC_PTR/ MDMAx_yy_CURR_DESC_PTR) 9-19
Current Address Register (DMAx_y_CURR_ADDR/ MDMAx_yy_CURR_ADDR) 9-20
Current Inner Loop Count Register (DMAx_y_CURR_X_COUNT/ MDMAx_yy_CURR_X_COUNT) 9-21
Current Outer Loop Count Register (DMAx_y_CURR_Y_COUNT/ MDMAx_yy_CURR_Y_COUNT) 9-22
Peripheral Map Register (DMAx_y_PERIPHERAL_MAP/ MDMAx_yy_PERIPHERAL_MAP) 9-23
Interrupt Status Register (DMAx_y_IRQ_STATUS/MDMAx_yy_IRQ_STATUS) 9-26
Flex Descriptor Structure 9-28
DMA Operation Flow 9-31
DMA Startup 9-33
DMA Refresh 9-35
To Stop DMA Transfers 9-37
Two-Dimensional DMA 9-38
More 2D DMA Examples 9-39

Memory DMA 9-40
MDMA Bandwidth 9-42
DMA Performance Optimization 9-43
Prioritization and Traffic Control 9-45
DMA Traffic Control Counter Period Register (DMAx_TC_PER) and DMA Traffic Control Counter Register (DMAx_TC_CNT) 9-47
MDMA Priority and Scheduling 9-50
Urgent DMA Transfers 9-51
Software Management of DMA 9-53
Synchronization of Software and DMA 9-53
Single-Buffer DMA Transfers 9-55
Continuous Transfers Using Autobuffering 9-56
Descriptor Structures 9-58
Descriptor Queue Management 9-59
Descriptor Queue Using Interrupts on Every Descriptor 9-60
Descriptor Queue Using Minimal Interrupts 9-61
DMA Errors (Aborts) 9-63

SPI COMPATIBLE PORT CONTROLLERS

Interface Signals 10-4
Serial Peripheral Interface Clock Signal (SCK) 10-4
Serial Peripheral Interface Slave Select Input Signal 10-4
Master Out Slave In (MOSI) 10-5

Master In Slave Out (MISO) 10-5
Interrupt Output 10-6
SPI Registers 10-7
SPI Baud Rate Register (SPI_BAUD) 10-7
SPI Control Register (SPI_CTL) 10-8
SPI Flag Register (SPI_FLG) 10-11
Slave Select Inputs 10-13
Use of FLS Bits in SPI_FLG for Multiple Slave SPI Systems 10-14
SPI Status Register (SPI_STAT) 10-15
SPI Transmit Data Buffer Register (SPI_TDBR) 10-17
SPI Receive Data Buffer Register (SPI_RDBR) 10-18
SPI RDBR Shadow Register (SPI_SHADOW) 10-18
Register Functions 10-19
SPI Transfer Protocols 10-20
SPI General Operation 10-22
Clock Signals 10-23
Master Mode Operation 10-24
Transfer Initiation From Master (Transfer Modes) 10-25
Slave Mode Operation 10-26
Slave Ready for a Transfer 10-27
Error Signals and Flags 10-27
Mode Fault Error (MODF) 10-27
Transmission Error (TXE) 10-29

Reception Error (RBSY) 10-29
Transmit Collision Error (TXCOL) 10-29
Beginning and Ending an SPI Transfer 10-29
DMA 10-31
DMA Functionality 10-31
Master Mode DMA Operation 10-32
Slave Mode DMA Operation 10-35
Timing 10-37

PARALLEL PERIPHERAL INTERFACE

PPI Memory-Mapped Registers 11-3
PPI Control Register (PPIx_CONTROL) 11-3
PPI Status Register (PPIx_STATUS) 11-9
Delay Count Register (PPIx_DELAY) 11-11
Transfer Count Register (PPIx_COUNT) 11-12
Lines Per Frame Register (PPIx_FRAME) 11-13
ITU-R 656 Modes 11-14
ITU-R 656 Background 11-14
ITU-R 656 Input Modes 11-18
Entire Field 11-19
Active Video Only 11-19
Vertical Blanking Interval (VBI) Only 11-19
ITU-R 656 Output Mode 11-20
Frame Synchronization in ITU-R 656 Modes 11-20

General-Purpose PPI Modes 11-21
Data Input (RX) Modes 11-24
No Frame Syncs 11-24
1, 2, or 3 External Frame Syncs 11-25
2 or 3 Internal Frame Syncs 11-26
Data Output (TX) Modes 11-26
No Frame Syncs 11-26
1 or 2 External Frame Syncs 11-27
1, 2, or 3 Internal Frame Syncs 11-28
Frame Synchronization in GP Modes 11-29
Modes With Internal Frame Syncs 11-29
Modes With External Frame Syncs 11-30
DMA Operation 11-31
Data Transfer Scenarios 11-33

SERIAL PORT CONTROLLERS

SPORT Operation 12-8
SPORT Disable 12-9
Setting SPORT Modes 12-10
Register Writes and Effective Latency 12-11
SPORTx_TCR1 and SPORTx_TCR2 Registers 12-11
SPORTx_RCR1 and SPORTx_RCR2 Registers 12-16
Data Word Formats 12-20
SPORTx_TX Register 12-21
SPORTx_RX Register 12-23

SPORTx_STAT Register 12-25
SPORT RX, TX, and Error Interrupts 12-27
PAB Errors 12-27
SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers 12-27
SPORTx_TFSDIV and SPORTx_RFSDIV Register 12-28
Clock and Frame Sync Frequencies 12-29
Maximum Clock Rate Restrictions 12-31
Frame Sync and Clock Example 12-31
Word Length 12-31
Bit Order 12-32
Data Type 12-32
Companding 12-33
Clock Signal Options 12-33
Frame Sync Options 12-34
Framed Versus Unframed 12-34
Internal Versus External Frame Syncs 12-36
Active Low Versus Active High Frame Syncs 12-37
Sampling Edge for Data and Frame Syncs 12-37
Early Versus Late Frame Syncs (Normal Versus Alternate Timing) 12-39
Data Independent Transmit Frame Sync 12-41
Moving Data Between SPORTs and Memory 12-42
Stereo Serial Operation 12-42

Multichannel Operation 12-46
SPORTx_MCMCn Registers 12-49
Multichannel Enable 12-50
Frame Syncs in Multichannel Mode 12-51
The Multichannel Frame 12-52
Multichannel Frame Delay 12-52
Window Size 12-53
Window Offset 12-54
SPORTx_CHNL Register 12-54
Other Multichannel Fields in SPORTx_MCMC2 12-55
Channel Selection Register 12-55
SPORTx_MRCSn Registers 12-57
SPORTx_MTCSn Registers 12-58
Multichannel DMA Data Packing 12-60
Support for H.100 Standard Protocol 12-61
2X Clock Recovery Control 12-61
SPORT Pin/Line Terminations 12-62
Timing Examples 12-62

UART PORT CONTROLLER

Serial Communications 13-2
UART Control and Status Registers 13-3
UART Line Control Register (UART_LCR) 13-4
Modem Control Register (UART_MCR) 13-5
UART Line Status Register (UART_LSR) 13-6

UART Transmit Holding Register (UART_THR) 13-7
UART Receive Buffer Register (UART_RBR) 13-8
UART Interrupt Enable Register (UART_IER) 13-9
UART Interrupt Identification Register (UART_IIR) 13-10
UART Divisor Latch Registers (UART_DLL, UART_DLH) 13-12
UART Scratch Register (UART_SCR) 13-14
UART Global Control Register (UART_GCTL) 13-14
Non-DMA Mode 13-15
DMA Mode 13-16
Mixing Modes 13-17
IrDA Support 13-18
IrDA Transmitter Description 13-18
IrDA Receiver Description 13-19

PROGRAMMABLE FLAGS

Programmable Flag Memory-Mapped Registers (MMRs) 14-8
Flag Direction Register (FIOn_DIR) 14-8
Flag Value Registers 14-10
Flag Data Register (FIOn_FLAG_D) 14-12
Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers 14-13
Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T) 14-19

Flag Interrupt Generation Flow 14-20
Flag Mask Interrupt A Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S, FIOn_MASKA_T) 14-22
Flag Mask Interrupt B Registers (FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T) 14-28
Flag Polarity Register (FIOn_POLAR) 14-33
Flag Interrupt Sensitivity Register (FIOn_EDGE) 14-35
Flag Set on Both Edges Register (FIOn_BOTH) 14-37
Flag Input Enable Register (FIOn_INEN) 14-39
Performance/Throughput 14-41

TIMERS

General-Purpose Timers 15-1
Timer Registers 15-4
Timer Enable (TMRSy_ENABLE) Registers 15-5
Timer Disable (TMRSy_DISABLE) Registers 15-5
Timer Status (TMRSy_STATUS) Registers 15-8
Timer Configuration (TIMERx_CONFIG) Registers 15-12
Timer Counter (TIMERx_COUNTER) Registers 15-14
Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers 15-15
Using the Timer 15-17
Pulse Width Modulation (PWM_OUT) Mode 15-20
Output Pad Disable 15-20

Single Pulse Generation 15-22
Pulse Width Modulation Waveform Generation 15-22
Stopping the Timer in PWM_OUT Mode 15-23
Externally clocked PWM_OUT 15-24
PULSE_HI Toggle Mode 15-25
Pulse Width Count and Capture (WDTH_CAP) Mode 15-30
Autobaud Mode 15-38
External Event (EXT_CLK) Mode 15-40
Using the Timers With the PPI 15-41
Interrupts 15-42
Illegal States 15-44
Summary 15-47
Core Timer 15-49
Core Timer Control Register (TCNTL) 15-50
Core Timer Count Register (TCOUNT) 15-51
Core Timer Period Register (TPERIOD) 15-52
Core Timer Scale Register (TSCALE) 15-53
Watchdog Timer 15-53
Watchdog Timer Operation 15-54
Watchdog Count Register (WDOG_CNT) 15-54
Watchdog Status Register (WDOG_STAT) 15-55
Watchdog Control Register (WDOG_CTL) 15-56

EXTERNAL BUS INTERFACE UNIT

Overview 16-1
Block Diagram 16-2
Internal Memory Interfaces 16-3
External Memory Interfaces 16-5
EBIU Programming Model 16-7
Error Detection 16-8
Asynchronous Memory Interface 16-8
Asynchronous Memory Address Decode 16-9
Asynchronous Memory Global Control Register (EBIU_AMGCTL) 16-9
Asynchronous Memory Bank Control Registers (EBIU_AMBCTL0, EBIU_AMBCTL1) 16-12
Avoiding Bus Contention 16-13
ARDY Input Control 16-13
Programmable Timing Characteristics 16-16
Asynchronous Reads 16-16
Asynchronous Writes 16-18
Adding Additional Wait States 16-19
Byte Enables 16-22
Accessing 8-Bit/16-Bit Devices (16-Bit Data Packing is Enabled) 16-22
Accessing 32-Bit Devices (16-Bit Data Packing is Disabled) 16-23

SDRAM Controller (SDC) 16-24
Introduction 16-24
Definition of Terms 16-25
Bank Activate Command 16-26
Burst Length 16-26
Burst Stop Command 16-26
Burst Type 16-26
CAS Latency (CL) 16-27
CBR (CAS Before RAS) Refresh or Auto-Refresh 16-27
DQM3—0 Data I/O Mask Function 16-27
Internal Bank 16-28
Mode Register 16-28
Page Size 16-29
Precharge Command 16-29
SDRAM Bank 16-29
Self-Refresh 16-29
t_{RAS} 16-30
t_{RP} 16-30
t_{RCD} 16-30
t_{WR} 16-31
t_{RC} 16-31
t_{RFC} 16-31

t_{XSR} 16-31
t_{RRD} 16-31
Example SDRAM System Block Diagrams 16-32
SDQM Operation 16-33
Executing a Parallel Refresh Command 16-33
SDRAM Memory Global Control Register (EBIU_SDGCTL) 16-34
Setting SDRAM Clock Enables (SCTLE and SCK1E) 16-38
Entering and Exiting Self-Refresh Mode (SRFS) 16-40
Setting SDRAM Buffering Timing Option (EBUFE) 16-41
Selecting the CAS Latency Value (CL) 16-42
Selecting the RAS to CAS Delay (TRCD) 16-43
Selecting the Bank Activate Command Delay (TRAS) 16-43
Selecting the Precharge Delay (TRP) 16-44
Selecting the Write to Precharge Delay (TWR) 16-45
SDRAM Memory Bank Control Register (EBIU_SDBCTL) 16-46
SDRAM Control Status Register (EBIU_SDSTAT) 16-49
SDRAM Refresh Rate Control Register (EBIU_SDRRC) 16-50
SDRAM External Bank Address Decode 16-52
SDRAM Address Mapping 16-55
16-Bit Wide SDRAM Address Muxing 16-56
Data Mask (SDQM3—0) Encodings 16-57
SDC Operation 16-58
SDC Configuration 16-59

SDC Commands 16-61
Precharge Command 16-62
Bank Activate Command 16-62
Load Mode Register Command 16-64
Read/Write Command 16-65
Auto-Refresh Command 16-65
Self-Refresh Command 16-65
No Operation/Command Inhibit Commands 16-67
SDRAM Timing Specifications 16-67
SDRAM Performance 16-68
Bus Request and Grant 16-68
Operation 16-69

SYSTEM DESIGN

Pin Descriptions 17-1
Recommendations for Unused Pins 17-1
Resetting the Processor 17-1
Booting the Processor 17-2
Managing Clocks 17-5
Managing Core and System Clocks 17-5
Configuring and Servicing Interrupts 17-6
Semaphores 17-6
Example Code for Query Semaphore 17-7
Data Delays, Latencies and Throughput 17-8
Bus Priorities 17-8

External Memory Design Issues 17-9
Example Asynchronous Memory Interfaces (16-Bit Memories) 17-9
Example Asynchronous Memory Interfaces (32-Bit Memories) 17-10
High Frequency Design Considerations 17-11
Point-to-Point Connections on Serial Ports 17-12
Signal Integrity 17-12
Decoupling Capacitors and Ground Planes 17-13
Oscilloscope Probes 17-14
Recommended Reading 17-15

BLACKFIN PROCESSOR DEBUG

Watchpoint Unit 18-1
Instruction Watchpoints 18-4
Instruction Watchpoint Address Registers (WPIAn) 18-5
Instruction Watchpoint Address Count Registers (WPIACNTn) 18-6
Instruction Watchpoint Address Control Register (WPIACTL) 18-7
Data Address Watchpoints 18-10
Data Watchpoint Address Registers (WPDAn) 18-11
Data Watchpoint Address Count Value Registers (WPDACNTn) 18-11
Data Watchpoint Address Control Register (WPDACTL) 18-12
Watchpoint Status Register (WPSTAT) 18-12

Trace Unit 18-14
Trace Buffer Control Register (TBUFCTL) 18-16
Trace Buffer Status Register (TBUFSTAT) 18-17
Trace Buffer Register (TBUF) 18-17
Code to Recreate the Execution Trace in Memory 18-18
Performance Monitoring Unit 18-19
Performance Monitor Counter Registers (PFCNTRn) 18-19
Performance Monitor Control Register (PFCTL) 18-20
Event Monitor Table 18-22
Cycle Counter 18-23
Execution Cycle Count Registers (CYCLES and CYCLES2) 18-24
Product Identification Registers 18-25
DSP Device ID Register (DSPID) 18-26

BLACKFIN PROCESSOR CORE MMR ASSIGNMENTS

L1 Data Memory Controller Registers A-2
L1 Instruction Memory Controller Registers A-5
Interrupt Controller Registers A-8
Core Timer Registers A-10
Debug, MP, and Emulation Unit Registers A-10
Trace Unit Registers A-11
Watchpoint and Patch Registers A-11
Performance Monitor Registers A-13

SYSTEM MMR ASSIGNMENTS

Dynamic Power Management Registers B-2
System Reset and Interrupt Control Registers B-2
Watchdog Timer Registers B-5
Parallel Peripheral Interface (PPI) Registers B-5
UART Controller Registers B-6
SPI Controller Registers B-7
Timer Registers B-8
Programmable Flag Registers B-12
SPORT0 Controller Registers B-17
SPORT1 Controller Registers B-19
DMA Control Registers B-20
DMA Traffic Control Registers B-23
External Bus Interface Unit Registers B-24

NUMERIC FORMATS

Unsigned or Signed: Two's-Complement Format C-1
Integer or Fractional C-2
Binary Multiplication C-4
Fractional Mode And Integer Mode C-5
Block Floating-Point Format C-6

GLOSSARY

INDEX

PREFACE

Thank you for purchasing and developing systems using an enhanced Blackfin® processor from Analog Devices, Inc.

Purpose of This Manual

The *ADSP-BF561 Blackfin Processor Hardware Reference* provides information about the architecture of the ADSP-BF561 Blackfin processor. The architectural descriptions cover functional blocks, buses, and ports.

For programming information, see the *Blackfin Processor Programming Reference.* For timing, electrical, and package specifications, see the *ADSP-BF561 Embedded Processor Data Sheet*.

Intended Audience

The primary audience for this manual is a programmer who is familiar with Analog Devices Blackfin processors. This manual assumes that the audience has a working knowledge of the appropriate processor architecture and instruction set. Programmers who are unfamiliar with Analog Devices processors can use this manual, but should supplement it with other texts, such as the appropriate programming reference manuals and data sheets, that describe their target architecture

Manual Contents

The manual consists of:

- Chapter 1, “Introduction”
  Provides a high level overview of the processor. Architectural descriptions include functional blocks, buses, and ports, including all features that they support.
- Chapter 2, “Computational Units”
  Describes the arithmetic/logic units (ALUs), multiplier/accumulator units (MACs), shifter, and the set of video ALUs. The chapter also discusses data formats, data types, and register files.
- Chapter 3, “Operating Modes and States”
  Describes the three operating modes of the processor: Emulation mode, Supervisor mode, and User mode. The chapter also describes Idle state and Reset state.
- Chapter 4, “Program Sequencer”
  Describes the operation of the program sequencer, which controls program flow by providing the address of the next instruction to be executed. The chapter also discusses loops, subroutines, jumps, interrupts, and exceptions.
- Chapter 5, “Data Address Generators”
  Describes the Data Address Generators (DAGs), addressing modes, how to modify DAG and Pointer registers, memory address alignment, and DAG instructions.
- Chapter 6, “Memory”
  Describes L1/L2 memories. In particular, details their memory architecture, memory model, memory transaction model, and memory-mapped registers (MMRs). Discusses the instruction, data, and scratchpad memories, which are part of the Blackfin processor core.

- Chapter 7, “Chip Bus Hierarchy”
  Describes on-chip buses, including how data moves through the system. The chapter also discusses the system memory map, major system components, and the system interconnects.
- Chapter 8, “Dynamic Power Management”
  Describes system reset and power-up configuration, system clocking and control, and power management.
- Chapter 9, “Direct Memory Access”
  Describes the peripheral DMA and Memory DMA controllers. The peripheral DMA section discusses direct, block data movements between a peripheral with DMA access and internal or external memory spaces.

  The Memory DMA section discusses memory-to-memory transfer capabilities among the processor memories: L1, L2, external synchronous, and external asynchronous memories.
- Chapter 10, “SPI Compatible Port Controllers”
  Describes the Serial Peripheral Interface (SPI) port that provides an I/O interface to a variety of SPI compatible peripheral devices.
- Chapter 11, “Parallel Peripheral Interface”
  Describes the Parallel Peripheral Interfaces (PPI0, PPI1) of the processor. Each PPI is a half-duplex, bidirectional port accommodating up to 16 bits of data and used for digital video and data converter applications.
- Chapter 12, “Serial Port Controllers”
  Describes the two independent, synchronous Serial Port Controllers (SPORT0 and SPORT1) that provide an I/O interface to a variety of serial peripheral devices.

- Chapter 13, “UART Port Controller”
  Describes the Universal Asynchronous Receiver/ Transmitter (UART) port, which converts data between serial and parallel formats and includes modem control and interrupt handling hardware.
- Chapter 14, “Programmable Flags”
  Describes the programmable flags, or general-purpose I/O pins in the processor, including how to configure the pins as inputs and outputs, and how to generate interrupts.
- Chapter 15, “Timers”
  Describes the twelve general-purpose timers that can be configured in any of three modes; the core timer that can generate periodic interrupts for a variety of timing functions; and the watchdog timer that can implement software watchdog functions, such as generating events to the Blackfin processor core.
- Chapter 16, “External Bus Interface Unit”
  Describes the External Bus Interface Unit of the processor. The chapter also discusses the asynchronous memory interface, the SDRAM controller (SDC), related registers, and SDC configuration and commands.
- Chapter 17, “System Design”
  Describes how to use the processor as part of an overall system. It includes information about interfacing the processor to external memory chips, bus timing and latency numbers, semaphores, and a discussion of the treatment of unused pins.
- Chapter 18, “Blackfin Processor Debug”
  Describes the Blackfin processor debug functionality, which can be used for software debugging and complements some services often found in an operating system.

- Appendix A, "Blackfin Processor Core MMR Assignments"
  Lists the core memory-mapped registers, their addresses, and cross-references to text.
- Appendix B, "System MMR Assignments"
  Lists the system memory-mapped registers, their addresses, and cross-references to text.
- Appendix C, "Numeric Formats"
  Describes various aspects of the 16-bit data format. The chapter also describes how to implement a block floating-point format in software.
- Glossary
  Contains definitions of terms used in this book, including acronyms.

This manual is a companion document to the *Blackfin Processor Programming Reference*.

What's New in This Manual

This is Revision 1.2 of the *ADSP-BF561 Blackfin Processor Hardware Reference*. Changes to this manual from Revision 1.1 include corrections of typographic errors and reported document errata and an updated *Preface*. Users should refer to the latest *ADSP-BF561 Embedded Processor Data Sheet* for the current specifications of the silicon.

Technical or Customer Support

You can reach Analog Devices, Inc. Customer Support in the following ways:

- Visit the Embedded Processing and DSP products Web site at
  `http://www.analog.com/processors/technical_support`
- E-mail tools questions to
  `processor.tools.support@analog.com`
- E-mail processor questions to
  `processor.support@analog.com (World wide support)`
  `processor.europe@analog.com (Europe support)`
  `processor.china@analog.com (China support)`
- Phone questions to **1-800-ANALOGD**
- Contact your Analog Devices, Inc. local sales office or authorized distributor
- Send questions by mail to:

```
Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
USA
```

Supported Processors

The name *Blackfin* refers to a family of 16-bit, embedded processors. VisualDSP++® currently supports the following Blackfin families:

ADSP-BF51x, ADSP-BF52x, ADSP-BF53x, ADSP-BF54x, and ADSP-BF56x

Product Information

Product information can be obtained from the Analog Devices Web site, VisualDSP++ online Help system, and a technical library CD.

Analog Devices Web Site

The Analog Devices Web site, www.analog.com, provides information about a broad range of products—analog integrated circuits, amplifiers, converters, and digital signal processors.

To access a complete technical library for each processor family, go to http://www.analog.com/processors/technical_library. The manuals selection opens a list of current manuals related to the product as well as a link to the previous revisions of the manuals. When locating your manual title, note a possible errata check mark next to the title that leads to the current correction report against the manual.

Also note, MyAnalog.com is a free feature of the Analog Devices Web site that allows customization of a Web page to display only the latest information about products you are interested in. You can choose to receive weekly e-mail notifications containing updates to the Web pages that meet your interests, including documentation errata against all manuals. MyAnalog.com provides access to books, application notes, data sheets, code examples, and more.

Visit MyAnalog.com to sign up. If you are a registered user, just log on. Your user name is your e-mail address.

VisualDSP++ Online Documentation

Online documentation comprises the VisualDSP++ Help system, software tools manuals, hardware tools manuals, processor manuals, Dinkum Abridged C++ library, and FLEXnet License Tools software

documentation. You can search easily across the entire VisualDSP++ documentation set for any topic of interest.

For easy printing, supplementary Portable Documentation Format (`.pdf`) files for all manuals are provided on the VisualDSP++ installation CD.

Each documentation file type is described as follows.

File	Description
`.chm`	Help system files and manuals in Microsoft help format
`.htm` or `.html`	Dinkum Abridged C++ library and FLEXnet License Tools software documentation. Viewing and printing the `.html` files requires a browser, such as Internet Explorer 6.0 (or higher).
`.pdf`	VisualDSP++ and processor manuals in PDF format. Viewing and printing the `.pdf` files requires a PDF reader, such as Adobe Acrobat Reader (4.0 or higher).

Technical Library CD

The technical library CD contains seminar materials, product highlights, a selection guide, and documentation files of processor manuals, VisualDSP++ software manuals, and hardware tools manuals for the following processor families: Blackfin, SHARC®, TigerSHARC®, ADSP-218x, and ADSP-219x.

To order the technical library CD, go to `http://www.analog.com/processors/technical_library`, navigate to the manuals page for your processor, click the request CD check mark, and fill out the order form.

Data sheets, which can be downloaded from the Analog Devices Web site, change rapidly, and therefore are not included on the technical library CD. Technical manuals change periodically. Check the Web site for the latest manual revisions and associated documentation errata.

Conventions

Text conventions used in this manual are identified and described as follows. Note that additional conventions, which apply only to specific chapters, may appear throughout this document.

Example	Description
SWRST Software Reset register	Register names appear in UPPERCASE and in a special typeface. The descriptive names of registers are in mixed case and regular typeface.
MOSI $\overline{\text{RESET}}$	Pin names appear in UPPERCASE and in the special typeface. Active low signals appear with an $\overline{\text{OVERBAR}}$.
AMS[3:0]	Register, bit, and pin names in the text may refer to groups of registers or pins. A colon between numbers within brackets indicates a range of registers or pins. For example, $\overline{\text{AMS[3:0]}}$ indicates $\overline{\text{AMS3}}$, $\overline{\text{AMS2}}$, $\overline{\text{AMS1}}$, and $\overline{\text{AMS0}}$.
TFSx	A lowercase x in a register name indicates a set of registers. For example, TFSx indicates TFS2, TFS1 and TFS0.
0x101F b#1010	The 0x prefix indicates a hexadecimal number or address The b# prefix indicates a binary number.
	Note: For correct operation, ... A Note: provides supplementary information on a related topic. In the online version of this book, the word **Note** appears instead of this symbol.
	Caution: Incorrect device operation may result if ... **Caution:** Device damage may result if ... A Caution: identifies conditions or inappropriate usage of the product that could lead to undesirable results or product damage. In the online version of this book, the word **Caution** appears instead of this symbol.

Example	Description
	Warning: Injury to device users may result if ... A Warning: identifies conditions or inappropriate usage of the product that could lead to conditions that are potentially hazardous for devices users. In the online version of this book, the word **Warning** appears instead of this symbol.
(See Table P-1.) Click Here	In the .pdf version of this book, blue cross references have a hypertext link to the item being referenced. Click on the blue words (Table, Figure, or section names) to jump to the location. In the online help versions of this book, the blue cross references are replaced by the words “Click Here”. Click on the “Click Here” words to jump to the referenced location.

Additional conventions, which apply only to specific chapters, may appear throughout this document.

Combining these text conventions leads to the following typical forms for register names.

Table P-1. Format of Register Names

Pattern	Description

`TIMERx_CONFIG`	The *x* refers to multiple instances of the peripheral.
	`TIMER0_CONFIG`
	`TIMER1_CONFIG`
	`TIMER2_CONFIG`

`ICPLB_DATAn`	The *n* refers to multiple registers within the same peripheral or within the same core component. `ICPLB_DATA15`

`SPORTx_TCRn`	The combination of x and n indicates multiple instances of the peripheral *and* multiple registers within the same peripheral.
	`SPORT0_TCR0`
	`SPORT1_TCR1`

Table P-1. Format of Register Names (Cont'd)

Pattern	Description
MDMAx_yy_CONFIG	The yy represents MemDMA Stream 0 or 1, either Destination or Source. The x represents the DMA controller number (either 0 or 1). MDMA0_D0_CONFIG MDMA0_S0_CONFIG MDMA0_D1_CONFIG MDMA0_S1_CONFIG
DMAx_y_START_ADDR	The y represents the DMA channel number. DMA0_5_START_ADDR DMA1_8_START_ADDR

Register Diagram Conventions

Register diagrams use the following conventions:

- The descriptive name of the register appears at the top, followed by the short form of the name in parentheses.
- If the register is read-only (RO), write-1-to-set (W1S), or write-1-to-clear (W1C), this information appears under the name. Read/write is the default and is not noted. Additional descriptive text may follow.
- If any bits in the register do not follow the overall read/write convention, this is noted in the bit description after the bit name.
- If a bit has a short name, the short name appears first in the bit description, followed by the long name in parentheses.
- The reset value appears in binary in the individual bits and in hexadecimal to the right of the register.

- Bits marked *x* have an unknown reset value. Consequently, the reset value of registers that contain such bits is undefined or dependent on pin values at reset.
- Shaded bits are reserved.

To ensure upward compatibility with future implementations, write back the value that is read for reserved bits in a register, unless otherwise specified.

Figure P-1 shows an example of these conventions.

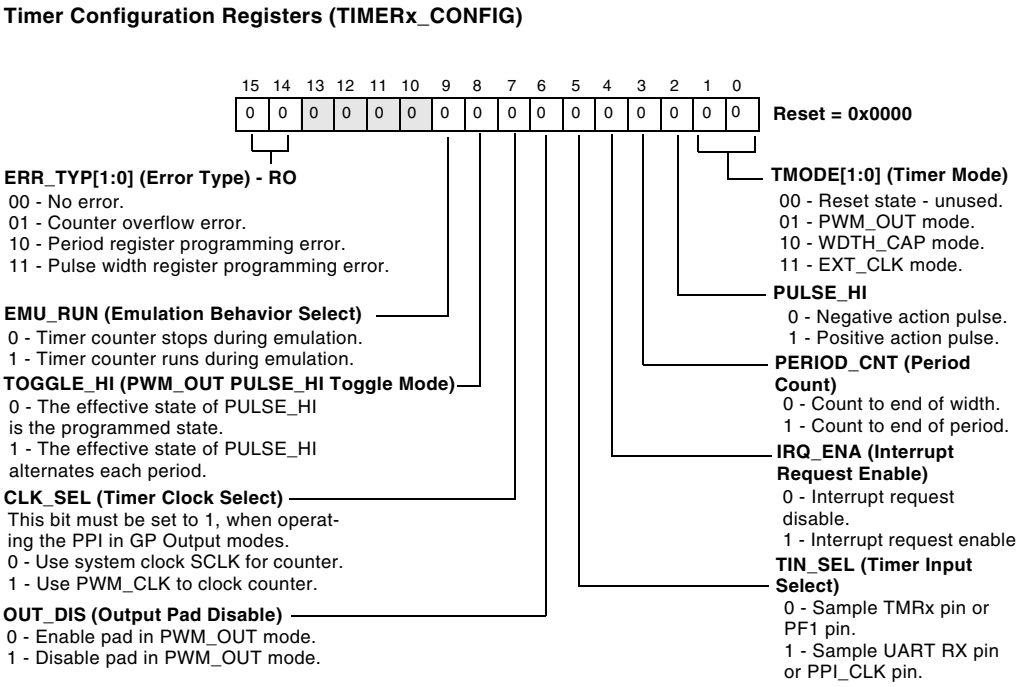

Figure P-1. Register Diagram Example

1 INTRODUCTION

The ADSP-BF561 is a member of the Blackfin processor family of products targeting consumer multimedia applications. At the heart of this device are two independent enhanced Blackfin processor cores that offer high performance and low power consumption while retaining their ease-of-use and code-compatibility benefits. The core architecture combines a dual-MAC signal processing engine, an orthogonal RISC-like microprocessor instruction set, flexible single-instruction multiple-data (SIMD) capabilities, and multimedia features into a single-instruction-set architecture.

Blackfin products feature Dynamic Power Management — the ability to vary both the voltage and frequency of operation, which optimizes the power consumption profile to the specific task.

ADSP-BF561 Peripherals

The ADSP-BF561 system peripherals include the following:

- Parallel Peripheral Interfaces (PPIs)
- Serial Ports (SPORTs)
- Serial Peripheral Interface (SPI)
- General-Purpose Timers
- Universal Asynchronous Receiver Transmitter (UART)

- Watchdog Timers
- General Purpose I/O (Programmable Flags)

These peripherals are connected to the core via several high bandwidth buses, as shown in Figure 1-1.

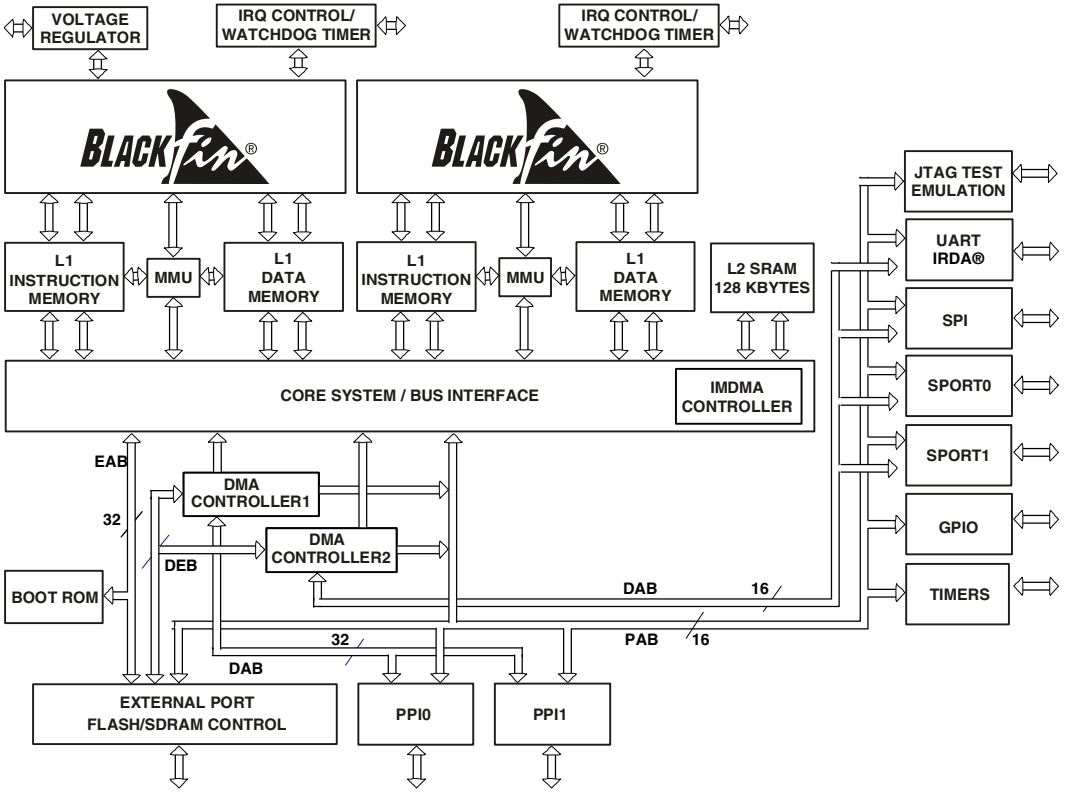

Figure 1-1. ADSP-BF561 Block Diagram

All of the peripherals, except for general-purpose I/O and Timers, are supported by a flexible DMA structure, including two DMA controllers (DMA1 and DMA2) and an internal memory DMA controller (IMDMA). In each of DMA1 and DMA2, there are twelve programmable

peripheral DMA channels and two separate memory DMA streams dedicated to data transfers between the DSP's memory spaces, which include external SDRAM and asynchronous memory. Multiple on-chip buses provide enough bandwidth to keep the processor core running even when there is also activity on all of the on-chip and external peripherals.

ADSP-BF561 Core Architecture

The ADSP-BF561 has two identical Blackfin cores and each core contains two 16-bit multipliers, two 40-bit accumulators, two 40-bit arithmetic logic units (ALUs), four 8-bit video ALUs, and a 40-bit shifter, along with the functional units of each core as shown in Figure 1-2. The computational units process 8-, 16-, or 32-bit data from the register file. All the functional units of each core are described below.

The compute register file contains eight 32-bit registers. When performing compute operations on 16-bit operand data, the register file operates as 16 independent 16-bit registers. All operands for compute operations come from the multiported register file and instruction constant fields.

Each MAC can perform a 16- by 16-bit multiply per cycle, with accumulation to a 40-bit result. Signed and unsigned formats, rounding, and saturation are supported.

The ALUs perform a traditional set of arithmetic and logical operations on 16-bit or 32-bit data. Many special instructions are included to accelerate various signal processing tasks. These include bit operations such as field extract and population count, modulo 2^{32} multiply, divide primitives, saturation and rounding, and sign/exponent detection. The set of video instructions includes byte alignment and packing operations, 16-bit and 8-bit adds with clipping, 8-bit average operations, and 8-bit subtract/absolute value/accumulate (SAA) operations. Also provided are the compare/select and vector search instructions. For some instructions, two 16-bit ALU operations can be performed simultaneously on register pairs

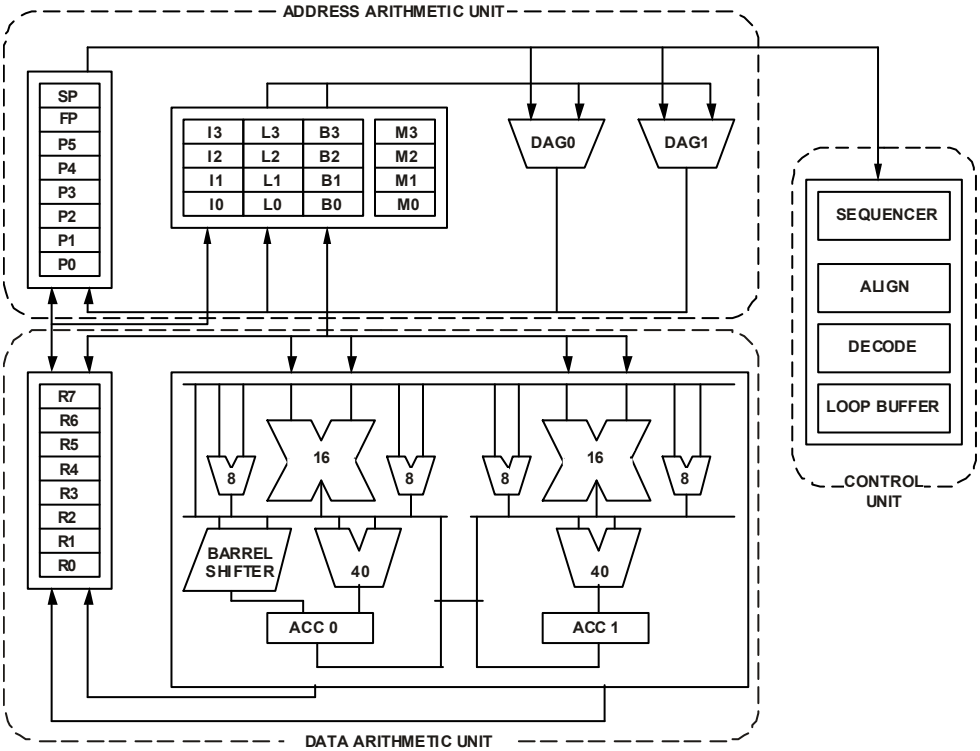

Figure 1-2. Core Architecture of Each ADSP-BF561 Core

(a 16-bit high half and 16-bit low half of a compute register). By also using the second ALU, quad 16-bit operations are possible.

The 40-bit shifter can deposit data and perform shifting, rotating, normalization, and extraction operations.

A program sequencer controls the instruction execution flow, including instruction alignment and decoding. For program flow control, the sequencer supports PC-relative and indirect conditional jumps (with static branch prediction), and subroutine calls. Hardware is provided to support zero-overhead looping. The architecture is fully interlocked, meaning that

there are no visible pipeline effects when executing instructions with data dependencies.

The address arithmetic unit provides two addresses for simultaneous dual fetches from memory. It contains a multiported register file consisting of four sets of 32-bit Index, Modify, Length, and Base registers (for circular buffering), and eight additional 32-bit pointer registers (for C-style indexed stack manipulation).

Blackfin products support a modified Harvard architecture in combination with a hierarchical memory structure. Level 1 (L1) memories typically operate at the full processor speed with little or no latency. At the L1 level, the instruction memory holds instructions only, the data memories hold data, and a dedicated scratchpad data memory stores stack and local variable information.

Multiple L1 memory blocks are provided, which may be configured as a mix of SRAM and cache. The Memory Management Unit (MMU) provides memory protection for individual tasks that may be operating on the core and can protect system registers from unintended access.

The ADSP-BF561 dual cores share an on-chip L2 memory system, which provides high speed SRAM access with somewhat longer latency than the L1 memory banks. The L2 memory is a unified instruction and data memory and can hold any mixture of code and data required by the system design.

The architecture provides three modes of operation: User, Supervisor, and Emulation. User mode has restricted access to a subset of system resources, thus providing a protected software environment. Supervisor and Emulation modes have unrestricted access to the system and core resources.

The Blackfin instruction set is optimized so that 16-bit op-codes represent the most frequently used instructions. Complex DSP instructions are encoded into 32-bit op-codes as multifunction instructions. Blackfin products support a limited multi-issue capability, where a 32-bit instruction can be issued in parallel with two 16-bit instructions. This

allows the programmer to use many of the core resources in a single instruction cycle.

The ADSP-BF561 assembly language uses an algebraic syntax. The architecture is also optimized for use with a C compiler.

Memory Architecture

The processor architecture structures memory as a single, unified 4 GB address space using 32-bit addresses. All resources, including internal memory, external memory, and I/O control registers, occupy separate sections of this common address space. The memory portions of this address space are arranged in a hierarchical structure to provide a good cost/performance balance of some very fast, low-latency on-chip memory as cache or SRAM, and larger, lower-cost and lower-performance off-chip memory systems. Figure 1-3 shows the memory map for the ADSP-BF561.

The memory DMA controller provides high-bandwidth data-movement capability. It can perform block transfers of code or data between the internal memory and the external memory spaces.

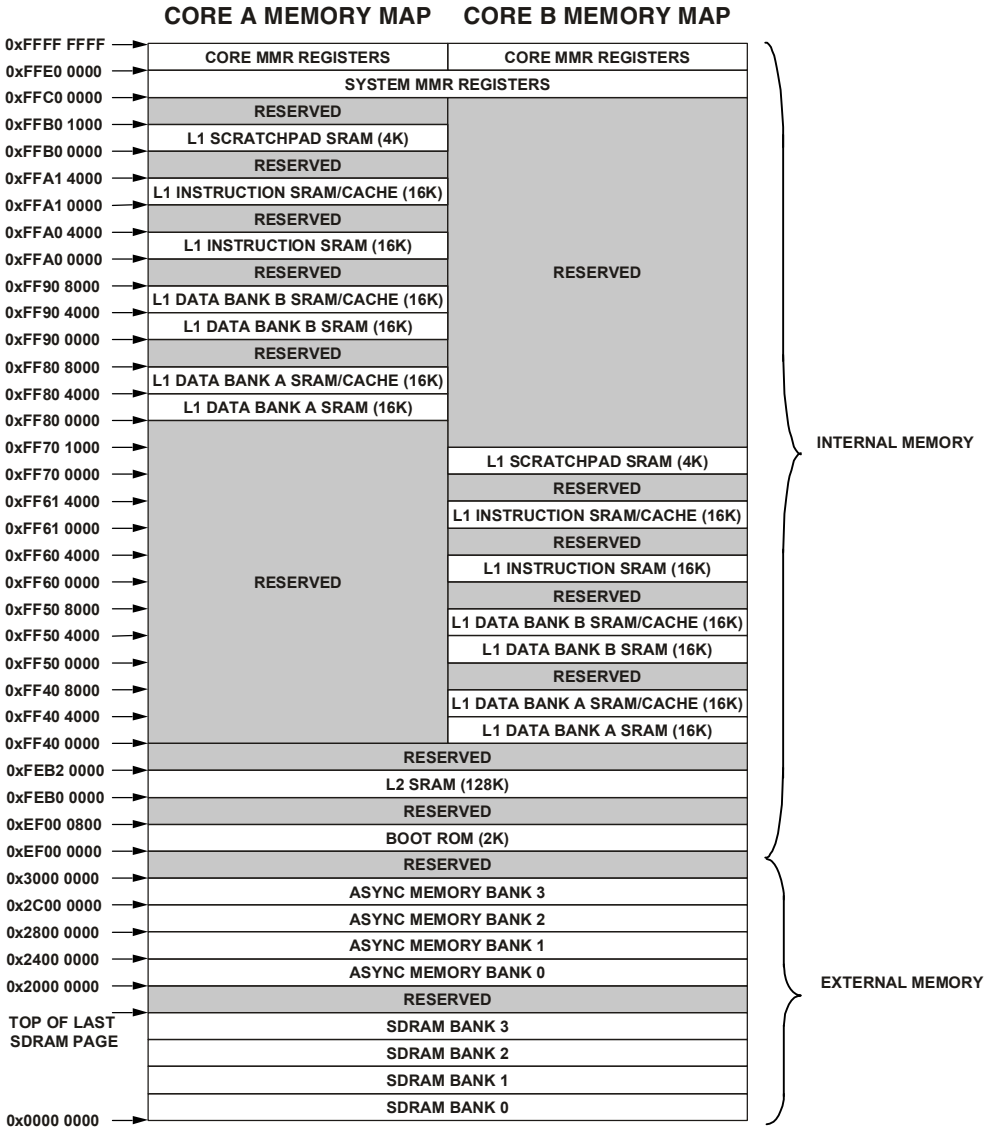

Figure 1-3. ADSP-BF561 Internal/External Memory Map

Internal Memory

Each core of the ADSP-BF561 has three blocks of on-chip memory that provide high-bandwidth access to the core:

- L1 instruction memory consisting of 32K bytes SRAM, of which 16K bytes can be configured as a four-way set-associative cache. This memory block is accessed at full processor speed.
- L1 data memory, consisting of two banks of 32K bytes each. One bank is always configured as SRAM while the other can be configured as SRAM or a two-way set associative cache.
- 4-Kbyte L1 scratchpad RAM, which runs at the full processor speed but is only accessible as data SRAM and cannot be configured as cache memory.

The ADSP-BF561 dual cores share a low latency, high-bandwidth on-chip Level 2 (L2) memory. It forms an on-chip memory hierarchy with L1 memory and provides much more capacity than L1 memory, but the latency is higher. The on-chip L2 memory is SRAM and can not be configured as cache. On-chip L2 memory is capable of storing both instructions and data.

External Memory

External (off-chip) memory is accessed via the External Bus Interface Unit (EBIU). This 32-bit interface provides a glueless connection to as many as four banks of synchronous DRAM (SDRAM) and as many as four banks of asynchronous memory devices including flash memory, EPROM, ROM, SRAM, and memory-mapped I/O devices.

The PC133-compliant SDRAM controller can be programmed to interface to up to 512 Mbytes of SDRAM.

The asynchronous memory controller can be programmed to control up to four banks of devices. Each bank occupies a 1 Mbyte segment regardless of the size of the devices used, so that these banks are only contiguous if each is fully populated with 1 Mbyte of memory.

I/O Memory Space

Blackfin products do not define a separate I/O space. All resources are mapped through the flat 32-bit address space. On-chip I/O devices have their control registers mapped into memory-mapped registers (MMRs) at addresses near the top of the 4 Gbyte address space. These are separated into two smaller blocks: one contains the control MMRs for all core functions and the other contains the registers needed for setup and control of the on-chip peripherals outside of the core. The MMRs are accessible only in Supervisor mode. They appear as reserved space to on-chip peripherals.

Event Handling

Each core on the ADSP-BF561 contains an event controller that handles all asynchronous and synchronous events to the core. ADSP-BF561 event handling supports both nesting and prioritization. Nesting allows multiple event service routines to be active simultaneously. Prioritization ensures that servicing a higher-priority event takes precedence over servicing a lower-priority event. The controller provides support for five different types of events:

- Emulation

  Emulation causes the processor to enter Emulation mode, allowing command and control of the processor via the JTAG interface.

- Reset

  Resets the processor.

- Nonmaskable Interrupt (NMI)

  The software watchdog timer or the NMI input signal to the processor generates this event. The NMI event is frequently used as a power-down indicator to initiate an orderly shutdown of the system.

- Exceptions

  Conditions such as data alignment violations and undefined instructions cause exceptions, which are synchronous to program flow. That is, the exception is taken before the instruction is allowed to complete.

- Interrupts

  Interrupts are asynchronous to program flow. These are caused by input pins, timers, and other peripherals.

Each event has an associated register to hold the return address and an associated return-from-event instruction. When an event is triggered, the state of the processor is saved on the supervisor stack.

Each core's event controller consists of two stages: the Core Event Controller (CEC) and the System Interrupt Controller (SIC). The CEC works with the SIC to prioritize and control all system events. Conceptually, interrupts from the peripherals arrive at the SIC and are routed directly into the general-purpose interrupts of the CEC.

Core Event Controller (CEC)

The CEC supports nine general-purpose interrupts (IVG15–7), in addition to the dedicated interrupt and exception events. Of these general-purpose interrupts, the two lowest-priority interrupts (IVG15–14) are recommended to be reserved for software interrupt handlers, leaving seven prioritized interrupt inputs to support peripherals.

System Interrupt Controller (SIC)

The System Interrupt Controller provides the mapping and routing of events from the many peripheral interrupt sources to the prioritized general-purpose interrupt inputs of the CEC. Although the ADSP-BF561 provides a default mapping, the user can alter the mappings and priorities of interrupt events by writing the appropriate values into the Interrupt Assignment Registers (IAR).

DMA Support

The ADSP-BF561 has multiple, independent DMA controllers that support automated data transfers with minimal overhead for the core. DMA transfers can occur between the ADSP-BF561 internal memories and any of its DMA-capable peripherals. Additionally, DMA transfers can be accomplished between any of the DMA-capable peripherals and external devices connected to the external memory interfaces, including the SDRAM controller and the asynchronous memory controller. DMA-capable peripherals include the SPORTs, SPI port, UART, and PPIs. Each individual DMA-capable peripheral has at least one dedicated DMA channel.

The DMA controller supports both 1-dimensional (1D) and 2-dimensional (2D) DMA transfers. DMA transfer initialization can be implemented from registers or from sets of parameters called descriptor blocks.

The 2D DMA capability supports arbitrary row and column sizes up to 64K elements by 64K elements, and arbitrary row and column step sizes up to +/- 32K elements. Furthermore, the column step size can be less than the row step size, allowing implementation of interleaved data streams. This feature is especially useful in video applications where data can be de-interleaved on the fly.

Examples of DMA types supported include:

- A single, linear buffer that stops upon completion
- A circular, auto-refreshing buffer that interrupts on each full or fractionally full buffer
- 1-D or 2-D DMA using a linked list of descriptors
- 2-D DMA using an array of descriptors specifying only the base DMA address within a common page

In addition to the dedicated peripheral DMA channels, there is a separate memory DMA channel provided for transfers between the various memories of the system. This enables transfers of blocks of data between any of the memories—including external SDRAM, ROM, SRAM, and flash memory—with minimal processor intervention. Memory DMA transfers can be controlled by a very flexible descriptor-based methodology or by a standard register-based autobuffer mechanism.

Another internal memory DMA (IMDMA) controller supports data transfer between internal memories only.

External Bus Interface Unit

The External Bus Interface Unit on the ADSP-BF561 interfaces with a wide variety of industry-standard memory devices. The controller consists of an SDRAM controller and an asynchronous memory controller.

PC133 SDRAM Controller

The SDRAM controller provides an interface to four banks of industry-standard SDRAM devices or DIMMs. Fully compliant with the PC133 SDRAM standard, each bank can be configured to contain between 16 and 128 Mbytes of memory.

A set of programmable timing parameters is available to configure SDRAM banks to support slower memory devices. The memory interface is 32-bits wide for minimum device count, lower system cost, and high performance.

Asynchronous Controller

The asynchronous memory controller provides a configurable interface for up to four separate banks of memory or I/O devices. Each bank can be independently programmed with different timing parameters. This allows connection to a wide variety of memory devices, including SRAM, ROM, and flash EPROM, as well as I/O devices that interface with standard memory control lines. Each bank occupies a 1-Mbyte window in the processor address space, but if not fully populated, these are not made contiguous by the memory controller. The banks are 32-bits wide, for interfacing to a range of memories and I/O devices.

Parallel Peripheral Interface

The processor provides two Parallel Peripheral Interfaces (PPI) that can connect directly to parallel A/D and D/A converters, ITU-R 601/656 video encoders and decoders, and other general-purpose peripherals. The PPI consists of a dedicated input clock pin, up to 3 frame synchronization pins, and up to 16 data pins. The input clock supports parallel data rates up to half the system clock rate.

In ITU-R 656 modes, the PPI receives and parses a data stream of 8-bit or 10-bit data elements. On-chip decode of embedded preamble control and synchronization information is supported.

Three distinct ITU-R 656 modes are supported:

- Active Video Only - The PPI does not read in any data between the End of Active Video (EAV) and Start of Active Video (SAV) preamble symbols, or any data present during the vertical blanking intervals. In this mode, the control byte sequences are not stored to memory; they are filtered by the PPI.
- Vertical Blanking Only - The PPI only transfers Vertical Blanking Interval (VBI) data, as well as horizontal blanking information and control byte sequences on VBI lines.
- Entire Field - The entire incoming bitstream is read in through the PPI. This includes active video, control preamble sequences, and ancillary data that may be embedded in horizontal and vertical blanking intervals.

Though not explicitly supported, ITU-R 656 output functionality can be achieved by setting up the entire frame structure (including active video, blanking, and control information) in memory and streaming the data out the PPI in a frame sync-less mode. The processor's 2D DMA features facilitate this transfer by allowing the static frame buffer (blanking and control codes) to be placed in memory once, and simply updating the active video information on a per-frame basis.

The general-purpose modes of the PPI are intended to suit a wide variety of data capture and transmission applications. The modes are divided into four main categories, each allowing up to 16 bits of data transfer per `PPI_CLK` cycle:

- Data Receive with Internally Generated Frame Syncs
- Data Receive with Externally Generated Frame Syncs
- Data Transmit with Internally Generated Frame Syncs
- Data Transmit with Externally Generated Frame Syncs

These modes support ADC/DAC connections, as well as video communication with hardware signalling. Many of the modes support more than one level of frame synchronization. If desired, a programmable delay can be inserted between assertion of a frame sync and reception/transmission of data.

Serial Ports (SPORTs)

The ADSP-BF561 incorporates two dual-channel synchronous serial ports (SPORT0 and SPORT1) for serial and multiprocessor communications. The SPORTs support the following features:

- Bidirectional, I^2S capable operation.

  Each SPORT has two sets of independent transmit and receive pins, enabling 4 bidirectional channels of I^2S stereo audio per SPORT.

- Buffered (8-deep) transmit and receive ports

  Each port has a data register for transferring data words to and from other system components and shift registers for shifting data in and out of the data registers.

- Clocking

  Each transmit and receive port can either use an external serial clock or can generate its own in a wide range of frequencies.

- Word length

  Each SPORT supports serial data words from 3 to 32 bits in length, transferred in most significant bit first or least significant bit first format.

- Framing

  Each transmit and receive port can run with or without frame sync signals for each data word. Frame sync signals can be generated internally or externally, active high or low, and with either of two pulse widths and early or late frame sync.

- Companding in hardware

  Each SPORT can perform A-law or µ-law companding according to ITU recommendation G.711. Companding can be selected on the transmit and/or receive channel of the SPORT without additional latencies.

- DMA operations with single-cycle overhead

  Each SPORT can automatically receive and transmit multiple buffers of memory data. The processor can link or chain sequences of DMA transfers between a SPORT and memory.

- Interrupts

  Each transmit and receive port generates an interrupt upon completing the transfer of a data word or after transferring an entire data buffer or buffers through DMA.

- Multichannel capability

  Each SPORT supports 128 channels out of a 1024-channel window and is compatible with the H.100, H.110, MVIP-90, and HMVIP standards.

Serial Peripheral Interface (SPI) Port

The ADSP-BF561 has a SPI-compatible port that enables the processor to communicate with multiple SPI-compatible devices.

The SPI interface uses three pins for transferring data: two data pins and a clock pin. An SPI chip select input pin lets other SPI devices select the processor, and seven SPI chip select output pins let the processor select other SPI devices. The SPI select pins are reconfigured Programmable Flag pins. Using these pins, the SPI port provides a full-duplex, synchronous serial interface, which supports both master and slave modes and multi-master environments.

The SPI port's baud rate and clock phase/polarities are programmable, and it has an integrated DMA controller, configurable to support either transmit or receive data streams. The SPI's DMA controller can only service unidirectional accesses at any given time.

During transfers, the SPI port simultaneously transmits and receives by serially shifting data in and out of its two serial data lines. The serial clock line synchronizes the shifting and sampling of data on the two serial data lines.

Timers

There are twelve general-purpose programmable timer units in the ADSP-BF561. Each timer has an external pin that can be configured either as a Pulse Width Modulator (PWM) or timer output, as an input to clock the timer, or as a mechanism for measuring pulse widths of external events. These timer units can be synchronized to an external clock input connected to the `EXTCLK` pin, an external clock input to the `PPI_CLK` pin, or to the internal `SCLK`.

The timer units can be used in conjunction with the UART to measure the width of the pulses in the data stream to provide an autobaud detect function for a serial channel.

The timers can generate interrupts to the processor core to provide periodic events for synchronization, either to the processor clock or to a count of external signals.

In addition to the twelve general-purpose programmable timers, A core timer is also provided for each core. This extra timer is clocked by the internal processor clock and is typically used as a system tick clock for generation of operating system periodic interrupts.

UART Port

The ADSP-BF561 provides a full-duplex Universal Asynchronous Receiver/Transmitter (UART) port, which is fully compatible with PC-standard UARTs. The UART port provides a simplified UART interface to other peripherals or hosts, providing full-duplex, DMA-supported, asynchronous transfers of serial data. The UART port includes support for 5 to 8 data bits; 1 or 2 stop bits; and none, even, or odd parity. The UART port supports two modes of operation:

- Programmed I/O (PIO)
- The processor sends or receives data by writing or reading I/O-mapped UART registers. The data is double-buffered on both transmit and receive.
- Direct Memory Access (DMA)
- The DMA controller transfers both transmit and receive data. This reduces the number and frequency of interrupts required to transfer data to and from memory. The UART has two dedicated DMA channels, one for transmit and one for receive. These DMA channels have lower priority than most DMA channels because of their relatively low service rates.

The UART port's baud rate, serial data format, error code generation and status, and interrupts can be programmed to support the following:

- Wide range of bit rates
- Data formats from 7 to 12 bits per frame
- Generation of maskable interrupts to the processor by both transmit and receive operations

In conjunction with the general-purpose timer functions, autobaud detection is supported.

The capabilities of the UART are further extended with support for the Infrared Data Association (IrDA®) Serial Infrared Physical Layer Link Specification (SIR) protocol.

Watchdog Timer

Each processor core in the ADSP-BF561 includes a 32-bit timer that can be used to implement a software watchdog function. A software watchdog can improve system availability by forcing the processor to a known state through generation of a hardware reset, non-maskable interrupt (NMI), or general purpose interrupt, if the timer expires before being reset by software. The programmer initializes the count value of the timer, enables the appropriate interrupt, then enables the timer. Thereafter, the software must reload the counter before it counts to zero from the programmed value. This protects the system from remaining in an unknown state where software, which would normally reset the timer, has stopped running due to an external noise condition or software error.

If configured to generate a hardware reset, the watchdog timer resets both the core and the ADSP-BF561 peripherals. After a reset, software can determine if the watchdog was the source of the hardware reset by interrogating a status bit in the watchdog timer control register.

The timer is clocked by the system clock (SCLK), at a maximum frequency of f_{SCLK}.

Programmable Flags

The ADSP-BF561 has 48 bidirectional programmable flag (PF) or general-purpose I/O pins, PF[47:0]. Each pin can be individually configured using the flag control, status, and interrupt registers.

- Flag Direction Control register – Specifies the direction of each individual PFx pin as input or output.
- Flag Control and Status registers – The ADSP-BF561 employs a "write one to modify" mechanism that allows any combination of individual flags to be modified in a single instruction, without affecting the level of any other flags. Four control registers are provided. One register is written in order to set flag values, one register is written in order to clear flag values, one register is written in order to toggle flag values, and one register is written in order to specify any number of flag values. Reading the Flag Status register allows software to interrogate the sense of the flags.
- Flag Interrupt Mask registers – The two Flag Interrupt Mask registers allow each individual PFx pin to function as an interrupt to the processor. Similar to the two Flag Control registers that are used to set and clear individual flag values, one Flag Interrupt Mask register sets bits to enable interrupt function, and the other Flag Interrupt Mask register clears bits to disable interrupt function. PFx pins defined as inputs can be configured to generate hardware interrupts, while output PFx pins can be triggered by software interrupts.
- Flag Interrupt Sensitivity registers – The two Flag Interrupt Sensitivity registers specify whether individual PFx pins are level- or edge-sensitive and specify—if edge-sensitive—whether just the

rising edge or both the rising and falling edges of the signal are significant. One register selects the type of sensitivity, and one register selects which edges are significant for edge-sensitivity.

Clock Signals

The processor can be clocked by an external crystal, a sine wave input, or a buffered, shaped clock derived from an external clock oscillator.

This external clock connects to the processor's CLKIN pin. The CLKIN input cannot be halted, changed, or operated below the specified frequency during normal operation. This clock signal should be a TTL-compatible signal.

The core clock (CCLK) and system peripheral clock (SCLK) are derived from the input clock (CLKIN) signal. An on-chip Phase Locked Loop (PLL) is capable of multiplying the CLKIN signal by a user-programmable (1x to 63x) multiplication factor (bounded by specified minimum and maximum VCO frequencies). The default multiplier is 10x, but it can be modified by a software instruction sequence. On-the-fly frequency changes can be made by simply writing to the PLL_DIV register.

All on-chip peripherals are clocked by the system clock (SCLK). The system clock frequency is programmable by means of the SSEL[3:0] bits of the PLL_DIV register.

Dynamic Power Management

The ADSP-BF561 provides four operating modes, each with a different performance/power profile. In addition, Dynamic Power Management provides the control functions to dynamically alter the processor core supply voltage to further reduce power dissipation. Control of clocking to each of the ADSP-BF561 peripherals also reduces power consumption.

Full On Operating Mode (Maximum Performance)

In the Full On mode, the phase-locked loop (PLL) is enabled, not bypassed, providing the maximum operational frequency. This is the normal execution state in which maximum performance can be achieved. The processor dual core and all enabled peripherals run at full speed.

Active Operating Mode (Moderate Power Savings)

In the Active mode, the PLL is enabled, but bypassed. Because the PLL is bypassed, the processor's core clocks (CCLK) and system clock (SCLK) run at the input clock (CLKIN) frequency. In this mode, the CLKIN to VCO multiplier ratio can be changed, although the changes are not realized until the Full On mode is entered. DMA access is available to appropriately configured L1 memories.

In the Active mode, it is possible to disable the PLL through the PLL Control register (PLL_CTL). If disabled, the PLL must be re-enabled before transitioning to the Full On or Sleep modes.

Sleep Operating Mode (High Power Savings)

The Sleep mode reduces power consumption by disabling the clock (CCLK) to the processor core. The PLL and system clock (SCLK), however, continue to operate in this mode. Typically an external event will wake up the processor. When in the Sleep mode, assertion of any interrupt will cause the processor to sense the value of the bypass bit (BYPASS) in the PLL control register (PLL_CTL). If bypass is disabled, the processor will transition to the Full On mode. If bypass is enabled, the processor will transition to the Active mode.

When in the Sleep mode, system DMA access to L1 memory is not supported.

Deep Sleep Operating Mode (Maximum Power Savings)

The Deep Sleep mode maximizes power savings by disabling the clocks to the processors dual cores (CCLK) and to all synchronous systems (SCLK). This powered-down mode can only be exited by assertion of the reset interrupt. When in Deep Sleep mode, assertion of the reset interrupt causes the processor to transition to the Full On mode or Active Mode, depending on the state of the BYPASS pin.

Hibernate State

For lowest possible power dissipation, this state allows the internal supply (V_{DDINT}) to be powered down, while keeping the I/O supply (V_{DDEXT}) running. Although not strictly an operating mode like the four modes detailed above, it is illustrative to view it as such in the diagram of Figure 8-6 on page 8-16. Since this feature is coupled to the on-chip switching regulator controller, it is discussed in detail in "Powering Down the Core (Hibernate State)" on page 8-28.

Voltage Regulation

The ADSP-BF561 provides an on-chip voltage regulator that can generate internal voltage levels (0.8 V to 1.2 V) from an external 2.25 V to 3.6 V supply. Figure 1-4 shows the typical external components required to complete the power management system. The regulator controls the internal logic voltage levels and is programmable with the Voltage Regulator Control Register (VR_CTL) in increments of 50 mV. To reduce standby power consumption, the internal voltage regulator can be programmed to remove power to the processor core while keeping I/O power supplied. While in this state, V_{DDEXT} can still be applied, eliminating the need for

external buffers. The regulator can also be disabled and bypassed at the user's discretion.

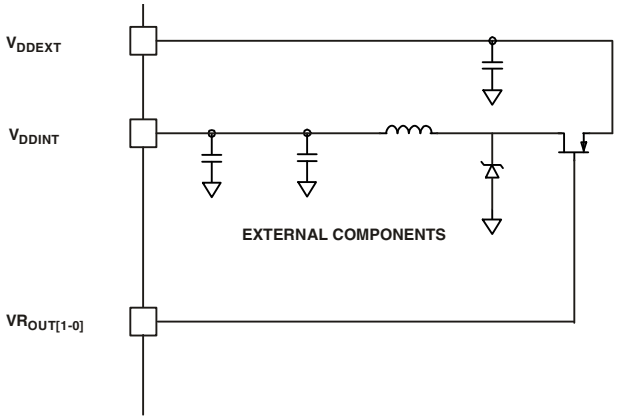

Figure 1-4. Voltage Regulator Circuit

Boot Modes

The ADSP-BF561 has three mechanisms for automatically loading internal L1 instruction memory after a reset.

- Boot from 8/16-bit external flash memory – The 8/16-bit flash boot routine located in boot ROM memory space is set up using Asynchronous Memory Bank 0. All configuration settings are set for the slowest device possible (3-cycle hold time; 15-cycle R/W access times; 4-cycle setup).
- Execute from 16-bit external memory – Execution starts from address 0x2000 0000 with 16-bit packing. The boot ROM is bypassed in this mode. All configuration settings are set for the slowest device possible (3-cycle hold time, 15-cycle read/write access times, 4-cycle setup).

- Boot from SPI serial EEPROM (16- or 24-bit addressable) – The SPI uses the PF2 output pin to select a single SPI EEPROM device, submits a read command at address 0x0000, and begins clocking data into the beginning of L1 instruction memory. A 16-bit addressable SPI-compatible EEPROM must be used.
- Boot from SPI host (slave mode) – A user-defined programmable flag pin is an output on the Blackfin processor and an input on the SPI host device. This flag allows the processor to hold off the host device from sending data during certain sections of the boot process. When this flag is de-asserted, the host can continue to send bytes to the processor.

For each of the boot modes, a boot loading protocol is used to transfer program and data blocks from an external memory device to their specified memory locations. Multiple memory blocks may be loaded by any boot sequence. Once all blocks are loaded, Core A program execution commences from the start of L1 instruction SRAM (0xFFA0 0000). Core B remains in a held-off state until a certain register bit is cleared. After that, Core B will start execution at address 0xFF60 0000.

In addition, bit 4 of the Reset Configuration Register can be set by application code to bypass the normal boot sequence during a software reset. For this case, the processor jumps directly to the beginning of L1 instruction memory.

Instruction Set Description

The Blackfin family assembly language instruction set employs an algebraic syntax designed for ease of coding and readability. The instructions have been specifically tuned to provide a flexible, densely encoded instruction set that compiles to a very small final memory size. The instruction set also provides fully featured multifunction instructions that allow the programmer to use many of the processor core resources in a single instruction. Coupled with many features more often seen on

microcontrollers, this instruction set is very efficient when compiling C and C++ source code. In addition, the architecture supports both user (algorithm/application code) and supervisor (O/S kernel, device drivers, debuggers, ISRs) modes of operation, allowing multiple levels of access to ADSP-BF561 resources.

The assembly language, which takes advantage of the processor's unique architecture, offers the following advantages:

- Seamlessly integrated DSP/CPU features are optimized for both 8-bit and 16-bit operations.
- A multi-issue load/store modified-Harvard architecture, which supports two 16-bit MAC or four 8-bit ALU plus two load/store operations plus two pointer updates per cycle.
- All registers, I/O, and memory are mapped into a unified 4G byte memory space, providing a simplified programming model.
- Microcontroller features, such as arbitrary bit and bit-field manipulation, insertion, and extraction; integer operations on 8-, 16-, and 32-bit data-types; and separate user and supervisor stack pointers.

Code density enhancements include intermixing of 16- and 32-bit instructions with no mode switching or code segregation. Frequently used instructions are encoded in 16 bits.

Development Tools

The ADSP-BF561 is supported with a complete set of CrossCore® software and hardware development tools, including Analog Devices emulators and the VisualDSP++ development environment. The same emulator hardware that supports other Analog Devices DSPs also fully emulates the ADSP-BF561.

The VisualDSP++ Integrated Development Environment lets programmers develop and debug an application. This environment includes an easy-to-use assembler that is based on an algebraic syntax, an archiver (librarian/library builder), a linker, a loader, a cycle-accurate instruction-level simulator, a C/C++ compiler, and a C/C++ run-time library that includes DSP and mathematical functions. A key point for these tools is C/C++ code efficiency. The compiler has been developed for efficient translation of C/C++ code to Blackfin assembly. The Blackfin processor has architectural features that improve the efficiency of compiled C/C++ code.

Debugging both C/C++ and assembly programs with the VisualDSP++ debugger, programmers can:

- View mixed C/C++ and assembly code (interleaved source and object information)
- Insert breakpoints
- Set conditional breakpoints on registers, memory, and stacks
- Trace instruction execution
- Perform linear or statistical profiling of program execution
- Fill, dump, and graphically plot the contents of memory
- Perform source level debugging
- Create custom debugger windows

The VisualDSP++ IDE lets programmers define and manage software development. Its dialog boxes and property pages let programmers configure and manage all development tools, including Color Syntax

Highlighting in the VisualDSP++ editor. These capabilities permit programmers to:

- Control how the development tools process inputs and generate outputs.
- Maintain a one-to-one correspondence with the tool's command line switches.

The VisualDSP++ Kernel (VDK) incorporates scheduling and resource management tailored specifically to address the memory and timing constraints of DSP programming. These capabilities enable engineers to develop code more effectively, eliminating the need to start from the very beginning, when developing new application code. The VDK features include threads, critical and unscheduled regions, semaphores, events, and device flags. The VDK also supports priority-based, pre-emptive, cooperative and time-sliced scheduling approaches. In addition, the VDK was designed to be scalable. If the application does not use a specific feature, the support code for that feature is excluded from the target system.

Because the VDK is a library, a developer can decide whether to use it or not. The VDK is integrated into the VisualDSP++ development environment but can also be used with standard command-line tools. The VDK development environment assists in managing system resources, automating the generation of various VDK-based objects, and visualizing the system state during application debug.

Analog Devices' DSP emulators use the IEEE 1149.1 JTAG test access port of the ADSP-BF561 to monitor and control the target board processor during emulation. The emulator provides full-speed emulation, allowing inspection and modification of memory, registers, and processor stacks. Non intrusive in-circuit emulation is assured by the use of the processor's JTAG interface—the emulator does not affect target system loading or timing.

In addition to the software and hardware development tools available from Analog Devices, third parties provide a wide range of tools

supporting the Blackfin DSP family. Hardware tools include the ADSP-BF561 EZ-KIT Lite® standalone evaluation/development cards. Third Party software tools include DSP libraries, real-time operating systems, and block diagram design tools.

2 COMPUTATIONAL UNITS

The processor contains two cores, each of which performs numeric processing for DSP and general control algorithms. The six computational units in each core are two arithmetic/logic units (ALUs), two multiplier/accumulator (multiplier) units, a shifter, and a set of video ALUs. These units get data from registers in the Data Register File. Computational instructions for these units provide fixed-point operations, and each computational instruction can execute every cycle.

The computational units handle different types of operations. The ALUs perform arithmetic and logic operations. The multipliers perform multiplication and execute multiply/add and multiply/subtract operations. The shifter executes logical shifts and arithmetic shifts and performs bit packing and extraction. The video ALUs perform Single-Instruction, Multiple-Data (SIMD) logical operations on specific 8-bit data operands.

Data moving in and out of the computational units goes through the Data Register File, which consists of eight registers, each 32 bits wide. In operations requiring 16-bit operands, the registers are paired, providing sixteen possible 16-bit registers.

The processor's assembly language provides access to the Data Register File. The syntax lets programs move data to and from these registers and specify a computation's data format at the same time.

Figure 2-1 provides a graphical guide to the other topics in this chapter. An examination of each computational unit provides details about its operation and is followed by a summary of computational instructions. Studying the details of the computational units, register files, and data buses leads to a better understanding of proper data flow for

computations. Next, details about the processor's advanced parallelism reveal how to take advantage of multifunction instructions.

Figure 2-1 shows the relationship between the Data Register File and the computational units—multipliers, ALUs, and shifter.

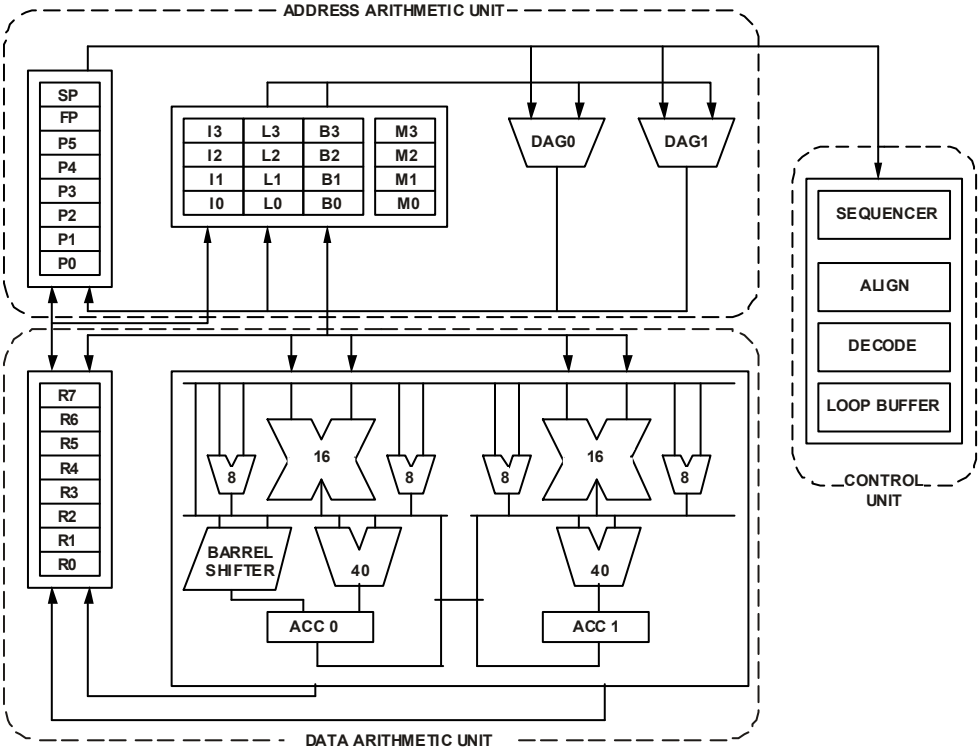

Figure 2-1. Processor Core Architecture

Single function multiplier, ALU, and shifter instructions have unrestricted access to the data registers in the Data Register File. Multifunction operations may have restrictions that are described in the section for that particular operation.

Two additional registers, A0 and A1, provide 40-bit accumulator results. These registers are dedicated to the ALUs and are used primarily for multiply-and-accumulate functions.

The traditional modes of arithmetic operations, such as fractional and integer, are specified directly in the instruction. Rounding modes are set from the ASTAT register, which also records status/conditions for the results of the computational operations.

Using Data Formats

Blackfin processors are primarily 16-bit, fixed-point machines. Most operations assume a two's-complement number representation, while others assume unsigned numbers or simple binary strings. Other instructions support 32-bit integer arithmetic, with further special features supporting 8-bit arithmetic and block floating point. For detailed information about each number format, see Appendix C, "Numeric Formats".

In the Blackfin processor family arithmetic, signed numbers are always in two's-complement format. These processors do not use signed-magnitude, one's-complement, binary coded decimal (BCD), or excess-n formats.

Binary String

The binary string format is the least complex binary notation; in it, 16 bits are treated as a bit pattern. Examples of computations using this format are the logical operations NOT, AND, OR, XOR. These ALU operations treat their operands as binary strings with no provision for sign bit or binary point placement.

Unsigned

Unsigned binary numbers may be thought of as positive and having nearly twice the magnitude of a signed number of the same length. The processor treats the least significant words of multiple precision numbers as unsigned numbers.

Signed Numbers: Two's-Complement

In Blackfin processor arithmetic, the word signed refers to two's-complement numbers. Most Blackfin processor family operations presume or support two's-complement arithmetic.

Fractional Representation: 1.15

Blackfin processor arithmetic is optimized for numerical values in a fractional binary format denoted by 1.15 ("one dot fifteen"). In the 1.15 format, 1 sign bit (the Most Significant Bit (MSB)) and 15 fractional bits represent values from –1 up to 0.999969.

Figure 2-2 shows the bit weighting for 1.15 numbers as well as some examples of 1.15 numbers and their decimal equivalents.

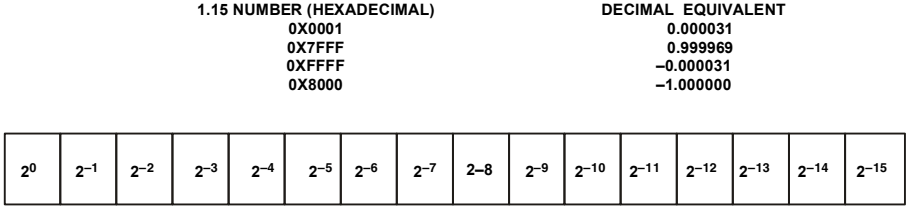

Figure 2-2. Bit Weighting for 1.15 Numbers

Register Files

The processor's computational units have three definitive register groups—a Data Register File, a Pointer Register File, and set of Data Address Generator (DAG) registers.

- The Data Register File receives operands from the data buses for the computational units and stores computational results.
- The Pointer Register File has pointers for addressing operations.
- The DAG registers are dedicated registers that manage zero-overhead circular buffers for DSP operations.

For more information, see Chapter 5, "Data Address Generators".

The processor register files appear in Figure 2-3.

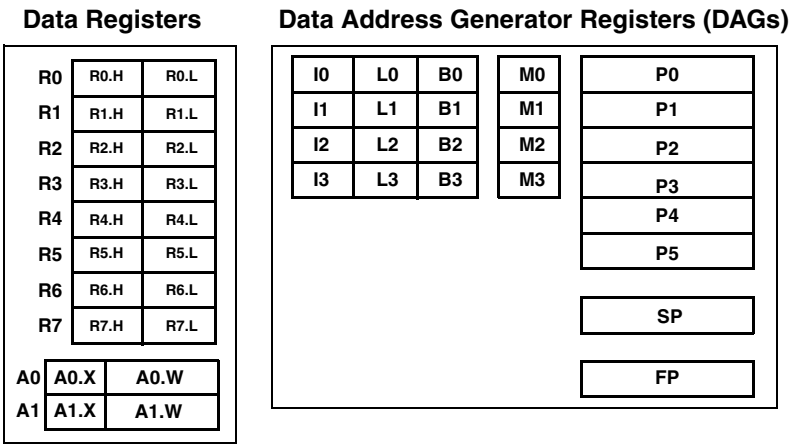

Figure 2-3. Register Files

In the processor, a word is 32 bits long; H denotes the high order 16 bits of a 32-bit register; L denotes the low order 16 bits of a 32-bit register. For example, `A0.W` contains the lower 32 bits of the 40-bit `A0` register; `A0.L` contains the lower 16 bits of `A0.W`, and `A0.H` contains the upper 16 bits of `A0.W`.

Data Register File

The Data Register File consists of eight registers, each 32 bits wide. Each register may be viewed as a pair of independent 16-bit registers. Each is denoted as the low half or high half. Thus the 32-bit register `R0` may be regarded as two independent register halves, `R0.L` and `R0.H`.

Three separate buses (two read, one write) connect the Register File to the L1 data memory, each bus being 32 bits wide. Transfers between the Data Register File and the data memory can move up to four 16-bit words of valid data in each cycle.

Accumulator Registers

In addition to the Data Register File, the processor has two dedicated, 40-bit accumulator registers. Each can be referred to as its 16-bit low half (`An.L`) or high half (`An.H`) plus its 8-bit extension (`An.X`). Each can also be referred to as a 32-bit register (`An.W`) consisting of the lower 32 bits, or as a complete 40-bit result register (`An`).

Pointer Register File

The general-purpose Address Pointer registers, also called P-registers, are organized as:

- 6-entry, P-register files `P[5:0]`
- Frame Pointers (`FP`) used to point to the current procedure's activation record

- Stack Pointer registers (SP) used to point to the last used location on the runtime stack. See mode dependent registers in Chapter 3, "Operating Modes and States".

P-registers are 32 bits wide. Although P-registers are primarily used for address calculations, they may also be used for general integer arithmetic with a limited set of arithmetic operations; for instance, to maintain counters. However, unlike the Data registers, P-register arithmetic does not affect the Arithmetic Status (ASTAT) register status flags.

DAG Register Set

DSP instructions primarily use the Data Address Generator (DAG) register set for addressing. The DAG register set consists of these registers:

- I[3:0] contain index addresses
- M[3:0] contain modify values
- B[3:0] contain base addresses
- L[3:0] contain length values

All DAG registers are 32 bits wide.

The I (Index) registers and B (Base) registers always contain addresses of 8-bit bytes in memory. The Index registers contain an effective address. The M (Modify) registers contain an offset value that is added to one of the Index registers or subtracted from it.

The B (Base) and L (Length) registers define circular buffers. The B register contains the starting address of a buffer, and the L register contains the length in bytes. Each L and B register pair is associated with the corresponding I register. For example, L0 and B0 are always associated with I0. However, any M register may be associated with any I register. For example, I0 may be modified by M3. For more information, see Chapter 5, "Data Address Generators".

Register File Instruction Summary

Table 2-1 lists the register file instructions. For more information about assembly language syntax, see the *Blackfin Processor Programming Reference*.

In Table 2-1, note the meaning of these symbols:

- Allreg denotes: `R[7:0]`, `P[5:0]`, `SP`, `FP`, `I[3:0]`, `M[3:0]`, `B[3:0]`, `L[3:0]`, `A0.X`, `A0.W`, `A1.X`, `A1.W`, `ASTAT`, `RETS`, `RETI`, `RETX`, `RETN`, `RETE`, `LC[1:0]`, `LT[1:0]`, `LB[1:0]`, `USP`, `SEQSTAT`, `SYSCFG`, `CYCLES`, and `CYCLES2`.
- A*n* denotes either ALU Result register `A0` or `A1`.
- Dreg denotes any Data Register File register.
- Sysreg denotes the system registers `ASTAT`, `SEQSTAT`, `SYSCFG`, `RETI`, `RETX`, `RETN`, `RETE`, or `RETS`, `LC[1:0]`, `LT[1:0]`, `LB[1:0]`, `CYCLES`, and `CYCLES2`.
- Preg denotes any Pointer register, `FP`, or `SP` register.
- Dreg_even denotes `R0`, `R2`, `R4`, or `R6`.
- Dreg_odd denotes `R1`, `R3`, `R5`, or `R7`.
- DPreg denotes any Data Register File register or any Pointer register, `FP`, or `SP` register.
- Dreg_lo denotes the lower 16 bits of any Data Register File register.
- Dreg_hi denotes the upper 16 bits of any Data Register File register.
- A*n*.L denotes the lower 16 bits of Accumulator `An.W`.
- A*n*.H denotes the upper 16 bits of Accumulator `An.W`.

- Dreg_byte denotes the low order 8 bits of each Data register.
- Option (X) denotes sign extended.
- Option (Z) denotes zero extended.
- * Indicates the flag may be set or cleared, depending on the result of the instruction.
- ** Indicates the flag is cleared.
- – Indicates no effect.

Table 2-1. Register File Instruction Summary

Instruction	ASTAT Status Flags						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AVS	AV1 AV1S	CC	V V_COPY VS
allreg = allreg ; [1]	–	–	–	–	–	–	–
A*n* = A*n* ;	–	–	–	–	–	–	–
A*n* = Dreg ;	–	–	–	–	–	–	–
Dreg_even = A0 ;	*	*	–	–	–	–	*
Dreg_odd = A1 ;	*	*	–	–	–	–	*
Dreg_even = A0, Dreg_odd = A1 ;	*	*	–	–	–	–	*
Dreg_odd = A1, Dreg_even = A0 ;	*	*	–	–	–	–	*
IF CC DPreg = DPreg ;	–	–	–	–	–	–	–
IF ! CC DPreg = DPreg ;	–	–	–	–	–	–	–
Dreg = Dreg_lo (Z) ;	*	**	**	–	–	–	**/–
Dreg = Dreg_lo (X) ;	*	*	**	–	–	–	**/–
A*n*.X = Dreg_lo ;	–	–	–	–	–	–	–
Dreg_lo = A*n*.X ;	–	–	–	–	–	–	–

Table 2-1. Register File Instruction Summary (Cont'd)

Instruction	ASTAT Status Flags						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AVS	AV1 AV1S	CC	V V_COPY VS
A*n*.L = Dreg_lo ;	–	–	–	–	–	–	–
A*n*.H = Dreg_hi ;	–	–	–	–	–	–	–
Dreg_lo = A0 ;	*	*	–	–	–	–	*
Dreg_hi = A1 ;	*	*	–	–	–	–	*
Dreg_hi = A1 ; Dreg_lo = A0 ;	*	*	–	–	–	–	*
Dreg_lo = A0 ; Dreg_hi = A1 ;	*	*	–	–	–	–	*
Dreg = Dreg_byte (Z) ;	*	**	**	–	–	–	**/–
Dreg = Dreg_byte (X) ;	*	*	**	–	–	–	**/–

1 Warning: Not all register combinations are allowed. For details, see the functional description of the Move Register instruction in the *Blackfin Processor Programming Reference*.

Data Types

The processor supports 32-bit words, 16-bit half words, and bytes. The 32- and 16-bit words can be integer or fractional, but bytes are always integers. Integer data types can be signed or unsigned, but fractional data types are always signed.

Table 2-2 illustrates the formats for data that resides in memory, in the register file, and in the accumulators. In the table, the letter *d* represents one bit, and the letter *s* represents one signed bit.

Some instructions manipulate data in the registers by sign-extending or zero-extending the data to 32 bits:

- Instructions zero-extend unsigned data
- Instructions sign-extend signed 16-bit half words and 8-bit bytes

Other instructions manipulate data as 32-bit numbers. In addition, two 16-bit half words or four 8-bit bytes can be manipulated as 32-bit values. For details, refer to the instructions in the *Blackfin Processor Programming Reference*.

In Table 2-2, note the meaning of these symbols:

- s = sign bit(s)
- d = data bit(s)
- "." = decimal point by convention; however, a decimal point does not literally appear in the number.
- Italics denotes data from a source other than adjacent bits.

Table 2-2. Data Formats

Representation in Memory	Representation in 32-bit Register
32.0 Unsigned Word dddd dddd dddd dddd dddd dddd dddd dddd	dddd dddd dddd dddd dddd dddd dddd dddd
32.0 Signed Word sddd dddd dddd dddd dddd dddd dddd dddd	sddd dddd dddd dddd dddd dddd dddd dddd
16.0 Unsigned Half Word dddd dddd dddd dddd	0000 0000 0000 0000 dddd dddd dddd dddd
16.0 Signed Half Word sddd dddd dddd dddd	ssss ssss ssss ssss sddd dddd dddd dddd

Table 2-2. Data Formats (Cont'd)

Representation in Memory	Representation in 32-bit Register
8.0 Unsigned Byte dddd dddd	0000 0000 0000 0000 0000 0000 dddd dddd
8.0 Signed Byte sddd dddd	ssss ssss ssss ssss ssss ssss sddd dddd
0.16 Unsigned Fraction .dddd dddd dddd dddd	0000 0000 0000 0000 .dddd dddd dddd dddd
1.15 Signed Fraction s.ddd dddd dddd dddd	ssss ssss ssss ssss s.ddd dddd dddd dddd
0.32 Unsigned Fraction .dddd dddd dddd dddd dddd dddd dddd dddd	.dddd dddd dddd dddd dddd dddd dddd dddd
1.31 Signed Fraction s.ddd dddd dddd dddd dddd dddd dddd dddd	s.ddd dddd dddd dddd dddd dddd dddd dddd
Packed 8.0 Unsigned Byte dddd dddd *dddd dddd* dddd dddd *dddd dddd*	dddd dddd *dddd dddd* dddd dddd *dddd dddd*
Packed 0.16 Unsigned Fraction .dddd dddd dddd dddd *.dddd dddd dddd dddd*	.dddd dddd dddd dddd *.dddd dddd dddd dddd*
Packed 1.15 Signed Fraction s.ddd dddd dddd dddd *s.ddd dddd dddd dddd*	s.ddd dddd dddd dddd *s.ddd dddd dddd dddd*

Endianess

Both internal and external memory are accessed in little endian byte order. For more information, see "Memory Transaction Model" on page 6-66.

ALU Data Types

Operations on each ALU treat operands and results as either 16- or 32-bit binary strings, except the signed division primitive (`DIVS`). ALU result

status bits treat the results as signed, indicating status with the overflow flags (AV0, AV1) and the negative flag (AN). Each ALU has its own sticky overflow flag, AV0S and AV1S. Once set, these bits remain set until cleared by writing directly to the ASTAT register. An additional V flag is set or cleared depending on the transfer of the result from both accumulators to the register file. Furthermore, the sticky VS bit is set with the V bit and remains set until cleared.

The logic of the overflow bits (V, VS, AV0, AV0S, AV1, AV1S) is based on two's-complement arithmetic. A bit or set of bits is set if the Most Significant Bit (MSB) changes in a manner not predicted by the signs of the operands and the nature of the operation. For example, adding two positive numbers must generate a positive result; a change in the sign bit signifies an overflow and sets AVn, the corresponding overflow flags. Adding a negative and a positive number may result in either a negative or positive result, but cannot cause an overflow.

The logic of the carry bits (AC0, AC1) is based on unsigned magnitude arithmetic. The bit is set if a carry is generated from bit 16 (the MSB). The carry bits (AC0, AC1) are most useful for the lower word portions of a multiword operation.

ALU results generate status information. For more information about using ALU status, see "ALU Instruction Summary" on page 2-28.

Multiplier Data Types

Each multiplier produces results that are binary strings. The inputs are interpreted according to the information given in the instruction itself (whether it is signed multiplied by signed, unsigned multiplied by unsigned, a mixture, or a rounding operation). The 32-bit result from the multipliers is assumed to be signed; it is sign-extended across the full 40-bit width of the A0 or A1 registers.

The processor supports two modes of format adjustment: the fractional mode for fractional operands (1.15 format with 1 sign bit and 15 fractional bits) and the integer mode for integer operands (16.0 format).

When the processor multiplies two 1.15 operands, the result is a 2.30 (2 sign bits and 30 fractional bits) number. In the fractional mode, the multiplier automatically shifts the multiplier product left one bit before transferring the result to the multiplier result register (`A0`, `A1`). This shift of the redundant sign bit causes the multiplier result to be in 1.31 format, which can be rounded to 1.15 format. The resulting format appears in Figure 2-4 on page 2-16.

In the integer mode, the left shift does not occur. For example, if the operands are in the 16.0 format, the 32-bit multiplier result would be in 32.0 format. A left shift is not needed and would change the numerical representation. This result format appears in Figure 2-5 on page 2-17.

Multiplier results generate status information when they update accumulators or when they are transferred to a destination register in the register file. For more information, see "Multiplier Instruction Summary" on page 2-39.

Shifter Data Types

Many operations in the shifter are explicitly geared to signed (two's-complement) or unsigned values—logical shifts assume unsigned magnitude or binary string values, and arithmetic shifts assume two's-complement values.

The exponent logic assumes two's-complement numbers. The exponent logic supports block floating point, which is also based on two's-complement fractions.

Shifter results generate status information. For more information about using shifter status, see "Shifter Instruction Summary" on page 2-52.

Arithmetic Formats Summary

Table 2-3, Table 2-4, Table 2-5, and Table 2-6 summarize some of the arithmetic characteristics of computational operations.

Table 2-3. ALU Arithmetic Formats

Operation	Operand Formats	Result Formats
Addition	Signed or unsigned	Interpret flags
Subtraction	Signed or unsigned	Interpret flags
Logical	Binary string	Same as operands
Division	Explicitly signed or unsigned	Same as operands

Table 2-4. Multiplier Fractional Modes Formats

Operation	Operand Formats	Result Formats
Multiplication	1.15 explicitly signed or unsigned	2.30 shifted to 1.31
Multiplication / Addition	1.15 explicitly signed or unsigned	2.30 shifted to 1.31
Multiplication / Subtraction	1.15 explicitly signed or unsigned	2.30 shifted to 1.31

Table 2-5. Multiplier Arithmetic Integer Modes Formats

Operation	Operand Formats	Result Formats
Multiplication	16.0 explicitly signed or unsigned	32.0 not shifted
Multiplication / Addition	16.0 explicitly signed or unsigned	32.0 not shifted
Multiplication / Subtraction	16.0 explicitly signed or unsigned	32.0 not shifted

Table 2-6. Shifter Arithmetic Formats

Operation	Operand Formats	Result Formats
Logical Shift	Unsigned binary string	Same as operands
Arithmetic Shift	Signed	Same as operands
Exponent Detect	Signed	Same as operands

Using Multiplier Integer and Fractional Formats

For multiply-and-accumulate functions, the processor provides two choices—fractional arithmetic for fractional numbers (1.15) and integer arithmetic for integers (16.0).

For fractional arithmetic, the 32-bit product output is format adjusted—sign-extended and shifted one bit to the left—before being added to accumulator `A0` or `A1`. For example, bit 31 of the product lines up with bit 32 of `A0` (which is bit 0 of `A0.X`), and bit 0 of the product lines up with bit 1 of `A0` (which is bit 1 of `A0.W`). The Least Significant Bit (LSB) is zero filled. The fractional multiplier result format appears in Figure 2-4.

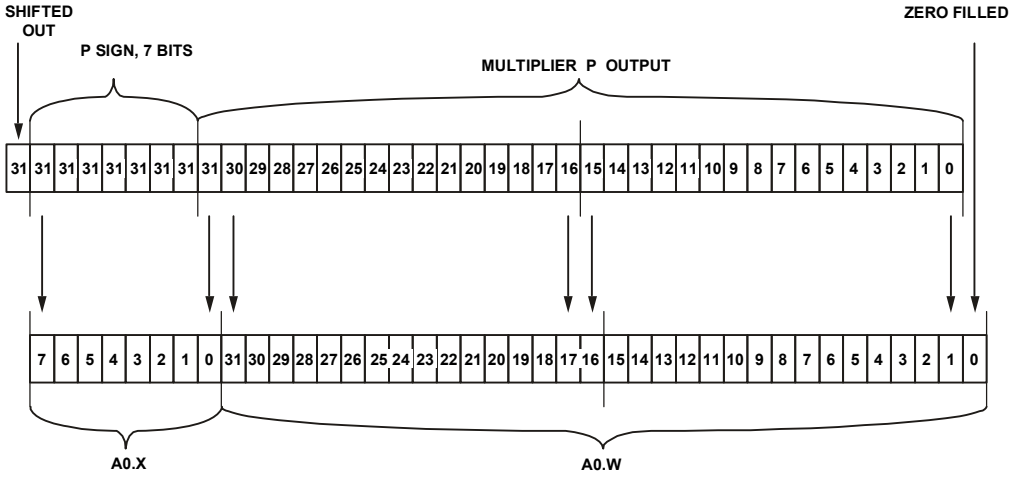

Figure 2-4. Fractional Multiplier Results Format

For integer arithmetic, the 32-bit product register is not shifted before being added to A0 or A1. Figure 2-5 shows the integer mode result placement.

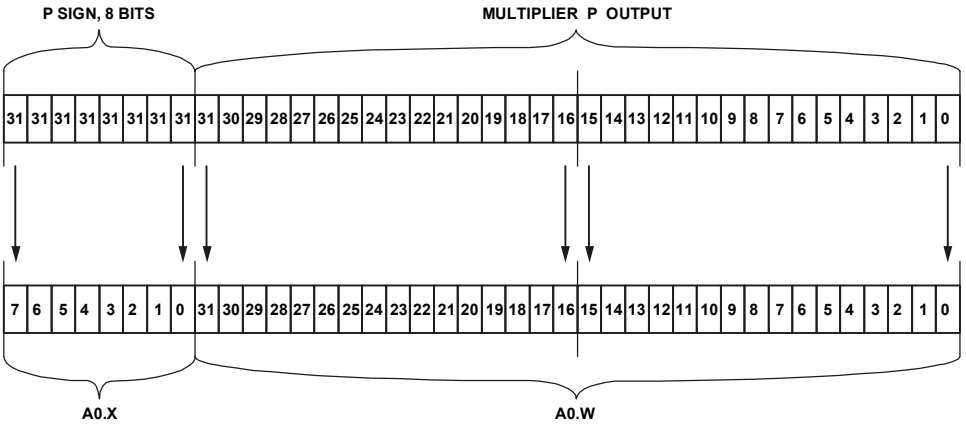

Figure 2-5. Integer Multiplier Results Format

With either fractional or integer operations, the multiplier output product is fed into a 40-bit adder/subtracter which adds or subtracts the new product with the current contents of the A0 or A1 register to produce the final 40-bit result.

Rounding Multiplier Results

On many multiplier operations, the processor supports multiplier results rounding (RND option). Rounding is a means of reducing the precision of a number by removing a lower order range of bits from that number's representation and possibly modifying the remaining portion of the number to more accurately represent its former value. For example, the original number will have N bits of precision, whereas the new number will have only M bits of precision (where N>M). The process of rounding, then, removes N – M bits of precision from the number.

The RND_MOD bit in the ASTAT register determines whether the RND option provides biased or unbiased rounding. For *unbiased* rounding, set RND_MOD bit = 0. For *biased* rounding, set RND_MOD bit = 1.

For most algorithms, unbiased rounding is preferred.

Unbiased Rounding

The *convergent* rounding method returns the number closest to the original. In cases where the original number lies exactly halfway between two numbers, this method returns the nearest even number, the one containing an LSB of 0. For example, when rounding the 3-bit, two's-complement fraction 0.25 (binary 0.01) to the nearest 2-bit, two's-complement fraction, the result would be 0.0, because that is the even-numbered choice of 0.5 and 0.0. Since it rounds up and down based on the surrounding values, this method is called *unbiased* rounding.

Unbiased rounding uses the ALU capability of rounding the 40-bit result at the boundary between bit 15 and bit 16. Rounding can be specified as part of the instruction code. When rounding is selected, the output register contains the rounded 16-bit result; the accumulator is never rounded.

The accumulator uses an unbiased rounding scheme. The conventional method of biased rounding adds a 1 into bit position 15 of the adder chain. This method causes a net positive bias because the midway value (when A0.L/A1.L = 0x8000) is always rounded upward.

The accumulator eliminates this bias by forcing bit 16 in the result output to 0 when it detects this midway point. Forcing bit 16 to 0 has the effect of rounding odd A0.L/A1.L values upward and even values downward, yielding a large sample bias of 0, assuming uniformly distributed values.

The following examples use *x* to represent any bit pattern (not all zeros). The example in Figure 2-6 shows a typical rounding operation for A0; the example also applies for A1.

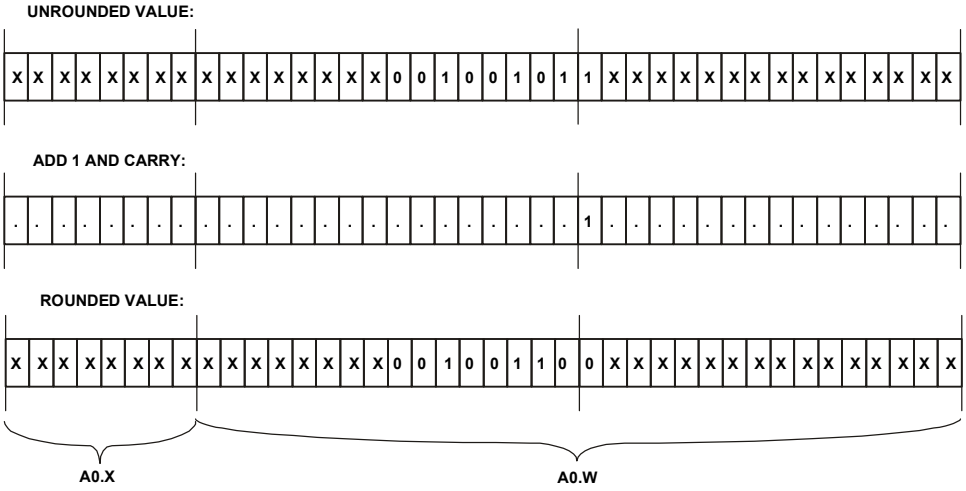

Figure 2-6. Typical Unbiased Multiplier Rounding

The compensation to avoid net bias becomes visible when all lower 15 bits are 0 and bit 15 is 1 (the midpoint value) as shown in Figure 2-7.

In Figure 2-7, A0 bit 16 is forced to 0. This algorithm is employed on every rounding operation, but is evident only when the bit patterns shown in the lower 16 bits of the next example are present.

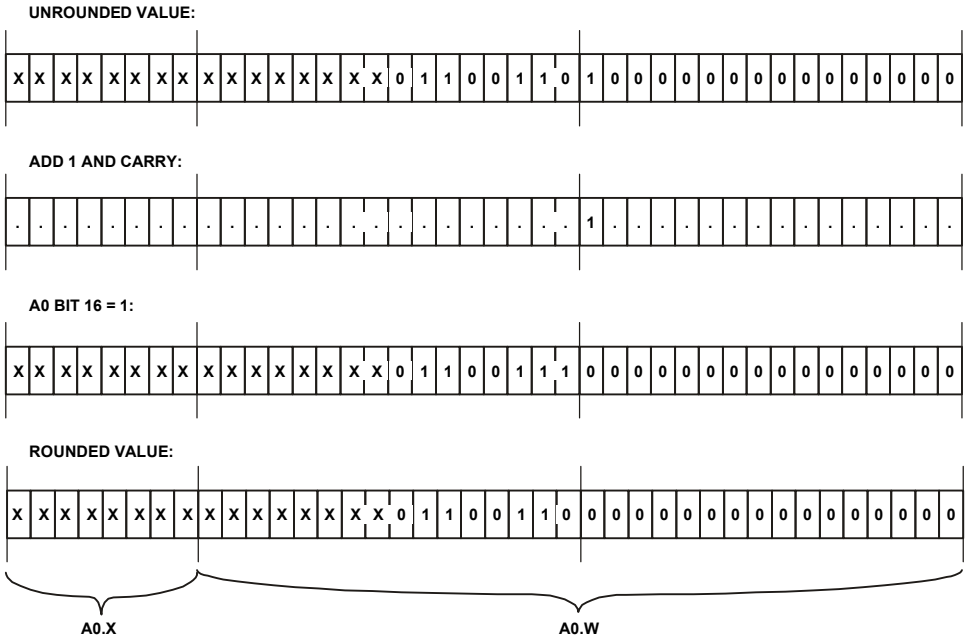

Figure 2-7. Avoiding Net Bias in Unbiased Multiplier Rounding

Biased Rounding

The *round-to-nearest* method also returns the number closest to the original. However, by convention, an original number lying exactly halfway between two numbers always rounds up to the larger of the two. For example, when rounding the 3-bit, two's-complement fraction 0.25 (binary 0.01) to the nearest 2-bit, two's-complement fraction, this method returns 0.5 (binary 0.1). The original fraction lies exactly midway between 0.5 and 0.0 (binary 0.0), so this method rounds up. Because it always rounds up, this method is called *biased* rounding.

The RND_MOD bit in the ASTAT register enables biased rounding. When the RND_MOD bit is cleared, the RND option in multiplier instructions uses the normal, unbiased rounding operation, as discussed in "Unbiased Rounding" on page 2-18.

When the RND_MOD bit is set (=1), the processor uses biased rounding instead of unbiased rounding. When operating in biased rounding mode, all rounding operations with A0.L/A1.L set to 0x8000 round up, rather than only rounding odd values up. For an example of biased rounding, see Table 2-7.

Table 2-7. Biased Rounding in Multiplier Operation

A0/A1 Before RND	Biased RND Result	Unbiased RND Result
0x00 0000 8000	0x00 0001 8000	0x00 0000 0000
0x00 0001 8000	0x00 0002 0000	0x00 0002 0000
0x00 0000 8001	0x00 0001 0001	0x00 0001 0001
0x00 0001 8001	0x00 0002 0001	0x00 0002 0001
0x00 0000 7FFF	0x00 0000 FFFF	0x00 0000 FFFF
0x00 0001 7FFF	0x00 0001 FFFF	0x00 0001 FFFF

Biased rounding affects the result only when the A0.L/A1.L register contains 0x8000; all other rounding operations work normally. This mode allows more efficient implementation of bit specified algorithms that use biased rounding (for example, the Global System for Mobile Communications (GSM) speech compression routines).

Truncation

Another common way to reduce the significant bits representing a number is to simply mask off the N – M lower bits. This process is known as *truncation* and results in a relatively large bias. Instructions that do not support rounding revert to truncation. The RND_MOD bit in ASTAT has no effect on truncation.

Special Rounding Instructions

The ALU provides the ability to round the arithmetic results directly into a data register with biased or unbiased rounding as described above. It also provides the ability to round on different bit boundaries. The options RND12, RND, and RND20 extract 16-bit values from bit 12, bit 16 and bit 20, respectively, and perform biased rounding regardless of the state of the RND_MOD bit in ASTAT.

For example:

```
R3.L = R4 (RND) ;
```

performs biased rounding at bit 16, depositing the result in a half word.

```
R3.L = R4 + R5 (RND12) ;
```

performs an addition of two 32-bit numbers, biased rounding at bit 12, depositing the result in a half word.

```
R3.L = R4 + R5 (RND20) ;
```

performs an addition of two 32-bit numbers, biased rounding at bit 20, depositing the result in a half word.

Using Computational Status

The multiplier, ALU, and shifter update the overflow and other status flags in the processor's Arithmetic Status (ASTAT) register. To use status conditions from computations in program sequencing, use conditional instructions to test the CC flag (bit 5, ASTAT register) after the instruction executes. This method permits monitoring each instruction's outcome. The ASTAT register is a 32-bit register, with some bits reserved. To ensure compatibility with future implementations, writes to this register should write back the values read from these reserved bits.

Arithmetic Status Register (ASTAT)

Figure 2-8 describes the ASTAT register. The processor updates the status bits in ASTAT, indicating the status of the most recent ALU, multiplier, or shifter operation.

Arithmetic Logic Unit (ALU)

The two ALUs perform arithmetic and logical operations on fixed-point data. ALU fixed-point instructions operate on 16-, 32-, and 40-bit fixed-point operands and output 16-, 32-, or 40-bit fixed-point results. ALU instructions include:

- Fixed-point addition and subtraction of registers
- Addition and subtraction of immediate values
- Accumulation and subtraction of multiplier results
- Logical AND, OR, NOT, XOR, bitwise XOR, Negate
- Functions: ABS, MAX, MIN, Round, division primitives

ALU Operations

Primary ALU operations occur on ALU0, while parallel operations occur on ALU1, which performs a subset of ALU0 operations.

Arithmetic Status Register (ASTAT)

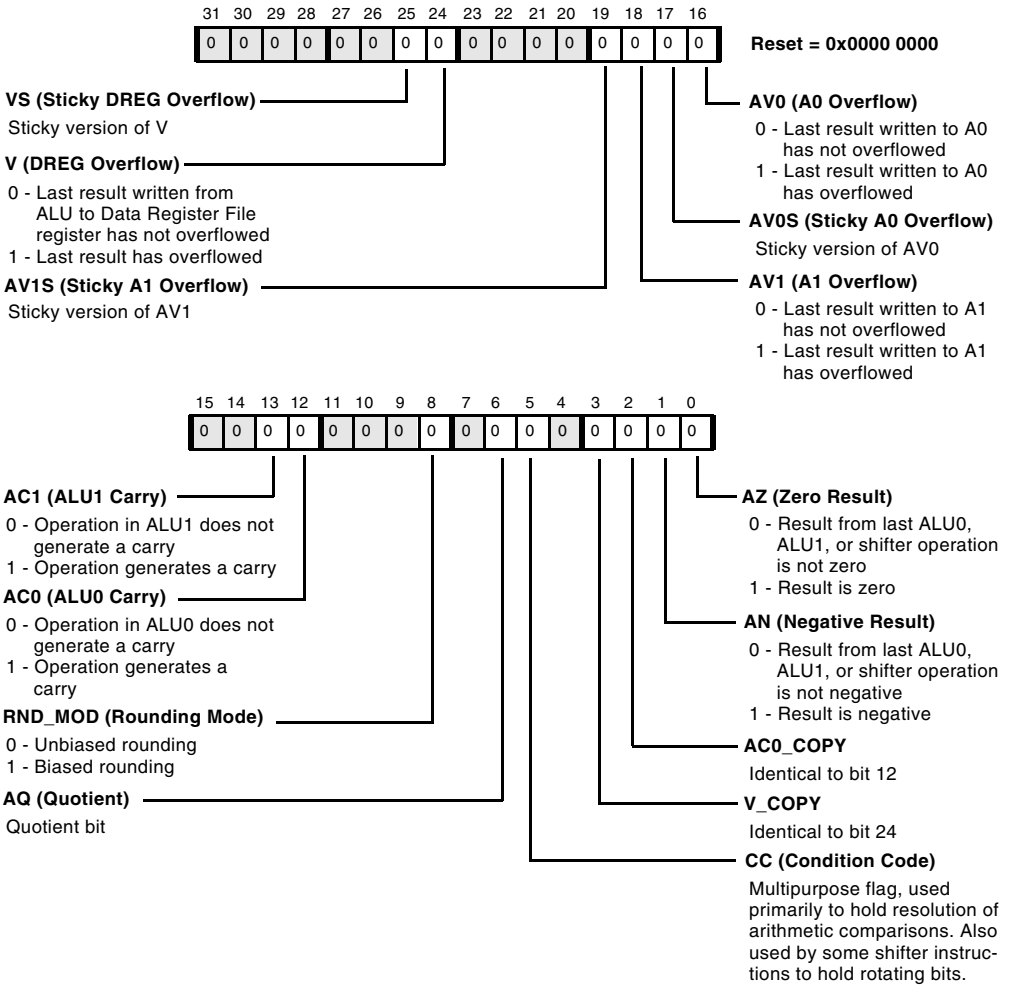

Figure 2-8. Arithmetic Status Register

Table 2-8 describes the possible inputs and outputs of each ALU.

Table 2-8. Inputs and Outputs of Each ALU

Input	Output
Two or four 16-bit operands	One or two 16-bit results
Two 32-bit operands	One 32-bit result
32-bit result from the multiplier	Combination of 32-bit result from the multiplier with a 40-bit accumulation result

Combining operations in both ALUs can result in four 16-bit results, two 32-bit results, or two 40-bit results generated in a single instruction.

Single 16-Bit Operations

In single 16-bit operations, any two 16-bit register halves may be used as the input to the ALU. An addition, subtraction, or logical operation produces a 16-bit result that is deposited into an arbitrary destination register half. ALU0 is used for this operation, because it is the primary resource for ALU operations.

For example:

```
R3.H = R1.H + R2.L (NS) ;
```

adds the 16-bit contents of `R1.H` (`R1` high half) to the contents of `R2.L` (`R2` low half) and deposits the result in `R3.H` (`R3` high half) with no saturation.

Dual 16-Bit Operations

In dual 16-bit operations, any two 32-bit registers may be used as the input to the ALU, considered as pairs of 16-bit operands. An addition, subtraction, or logical operation produces two 16-bit results that are deposited into an arbitrary 32-bit destination register. ALU0 is used for this operation, because it is the primary resource for ALU operations.

For example:

```
R3 = R1 +|- R2 (S) ;
```

adds the 16-bit contents of `R2.H` (`R2` high half) to the contents of `R1.H` (`R1` high half) and deposits the result in `R3.H` (`R3` high half) with saturation.

The instruction also subtracts the 16-bit contents of `R2.L` (`R2` low half) from the contents of `R1.L` (`R1` low half) and deposits the result in `R3.L` (`R3` low half) with saturation (see Figure 2-10 on page 2-34).

Quad 16-Bit Operations

In quad 16-bit operations, any two 32-bit registers may be used as the inputs to ALU0 and ALU1, considered as pairs of 16-bit operands. A small number of addition or subtraction operations produces four 16-bit results that are deposited into two arbitrary, 32-bit destination registers. Both ALU0 and ALU1 are used for this operation. Because there are only two 32-bit data paths from the Data Register File to the arithmetic units, the same two pairs of 16-bit inputs are presented to ALU1 as to ALU0. The instruction construct is identical to that of a dual 16-bit operation, and input operands must be the same for both ALUs.

For example:

```
R3 = R0 +|+ R1, R2 = R0 -|- R1 (S) ;
```

performs four operations:

- Adds the 16-bit contents of `R1.H` (`R1` high half) to the 16-bit contents of `R0.H (R0` high half) and deposits the result in `R3.H` with saturation.
- Adds `R1.L` to `R0.L` and deposits the result in `R3.L` with saturation.

- Subtracts the 16-bit contents of R1.H (R1 high half) from the 16-bit contents of the R0.H (R0 high half) and deposits the result in R2.H with saturation.
- Subtracts R1.L from R0.L and deposits the result in R2.L with saturation.

Explicitly, the four equivalent instructions are:

```
R3.H = R0.H + R1.H (S) ;
R3.L = R0.L + R1.L (S) ;
R2.H = R0.H - R1.H (S) ;
R2.L = R0.L - R1.L (S) ;
```

Single 32-Bit Operations

In single 32-bit operations, any two 32-bit registers may be used as the input to the ALU, considered as 32-bit operands. An addition, subtraction, or logical operation produces a 32-bit result that is deposited into an arbitrary 32-bit destination register. ALU0 is used for this operation, because it is the primary resource for ALU operations.

In addition to the 32-bit input operands coming from the Data Register File, operands may be sourced and deposited into the Pointer Register File, consisting of the eight registers P[5:0], SP, FP.

Instructions may not intermingle Pointer registers with Data registers.

For example:

```
R3 = R1 + R2 (NS) ;
```

adds the 32-bit contents of R2 to the 32-bit contents of R1 and deposits the result in R3 with no saturation.

```
R3 = R1 + R2 (S) ;
```

adds the 32-bit contents of R1 to the 32-bit contents of R2 and deposits the result in R3 with saturation.

Dual 32-Bit Operations

In dual 32-bit operations, any two 32-bit registers may be used as the input to ALU0 and ALU1, considered as a pair of 32-bit operands. An addition or subtraction produces two 32-bit results that are deposited into two 32-bit destination registers. Both ALU0 and ALU1 are used for this operation. Because only two 32-bit data paths go from the Data Register File to the arithmetic units, the same two 32-bit input registers are presented to ALU0 and ALU1.

For example:

```
R3 = R1 + R2, R4 = R1 - R2 (NS) ;
```

adds the 32-bit contents of R2 to the 32-bit contents of R1 and deposits the result in R3 with no saturation.

The instruction also subtracts the 32-bit contents of R2 from that of R1 and deposits the result in R4 with no saturation.

A specialized form of this instruction uses the ALU 40-bit result registers as input operands, creating the sum and differences of the A0 and A1 registers.

For example:

```
R3 = A0 + A1, R4 = A0 - A1 (S) ;
```

transfers to the result registers two 32-bit, saturated, sum and difference values of the ALU registers.

ALU Instruction Summary

Table 2-9 lists the ALU instructions. For more information about assembly language syntax and the effect of ALU instructions on the status flags, see the *Blackfin Processor Programming Reference*.

In Table 2-9, note the meaning of these symbols:

- Dreg denotes any Data Register File register.
- Preg denotes any Pointer register, `FP`, or `SP` register.
- Dreg_lo_hi denotes any 16-bit register half in any Data Register File register.
- Dreg_lo denotes the lower 16 bits of any Data Register File register.
- imm7 denotes a signed, 7-bit wide, immediate value.
- A*n* denotes either ALU Result register `A0` or `A1`.
- DIVS denotes a Divide Sign primitive.
- DIVQ denotes a Divide Quotient primitive.
- MAX denotes the maximum, or most positive, value of the source registers.
- MIN denotes the minimum value of the source registers.
- ABS denotes the absolute value of the upper and lower halves of a single 32-bit register.
- RND denotes rounding a half word.
- RND12 denotes saturating the result of an addition or subtraction and rounding the result on bit 12.
- RND20 denotes saturating the result of an addition or subtraction and rounding the result on bit 20.
- SIGNBITS denotes the number of sign bits in a number, minus one.

- EXPADJ denotes the lesser of the number of sign bits in a number minus one, and a threshold value.
- * Indicates the flag may be set or cleared, depending on the results of the instruction.
- ** Indicates the flag is cleared.
- – Indicates no effect.
- *d* indicates AQ contains the dividend MSB Exclusive-OR divisor MSB.

Table 2-9. ALU Instruction Summary

Instruction	ASTAT Status Flags						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	V V_COPY VS	AQ
Preg = Preg + Preg ;	–	–	–	–	–	–	–
Preg += Preg ;	–	–	–	–	–	–	–
Preg –= Preg ;	–	–	–	–	–	–	–
Dreg = Dreg + Dreg ;	*	*	*	–	–	*	–
Dreg = Dreg – Dreg (S) ;	*	*	*	–	–	*	–
Dreg = Dreg + Dreg, Dreg = Dreg – Dreg ;	*	*	*	–	–	*	–
Dreg_lo_hi = Dreg_lo_hi + Dreg_lo_hi ;	*	*	*	–	–	*	–
Dreg_lo_hi = Dreg_lo_hi – Dreg_lo_hi (S) ;	*	*	*	–	–	*	–
Dreg = Dreg +|+ Dreg ;	*	*	*	–	–	*	–
Dreg = Dreg +|– Dreg ;	*	*	*	–	–	*	–
Dreg = Dreg –|+ Dreg ;	*	*	*	–	–	*	–
Dreg = Dreg –|– Dreg ;	*	*	*	–	–	*	–
Dreg = Dreg +|+Dreg, Dreg = Dreg –|– Dreg ;	*	*	–	–	–	*	–

Table 2-9. ALU Instruction Summary (Cont'd)

Instruction	ASTAT Status Flags						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	V V_COPY VS	AQ
Dreg = Dreg +|– Dreg, Dreg = Dreg –|+ Dreg ;	*	*	–	–	–	*	–
Dreg = A*n* + A*n*, Dreg = A*n* – A*n* ;	*	*	*	–	–	*	–
Dreg += imm7 ;	*	*	*	–	–	*	–
Preg += imm7 ;	–	–	–	–	–	–	–
Dreg = (A0 += A1) ;	*	*	*	*	–	*	–
Dreg_lo_hi = (A0 += A1);	*	*	*	*	–	*	–
A0 += A1 ;	*	*	*	*	–	–	–
A0 –= A1 ;	*	*	*	*	–	–	–
DIVS (Dreg, Dreg) ;	*	*	*	*	–	–	*d*
DIVQ (Dreg, Dreg) ;	*	*	*	*	–	–	*d*
Dreg = MAX (Dreg, Dreg) (V) ;	*	*	–	–	–	**/–	–
Dreg = MIN (Dreg, Dreg) (V) ;	*	*	–	–	–	**/–	–
Dreg = ABS Dreg (V) ;	*	**	–	–	–	*	–
A*n* = ABS A*n* ;	*	**	–	*	*	*	–
A*n* = ABS A*n*, A*n* = ABS A*n* ;	*	**	–	*	*	*	–
A*n* = –A*n* ;	*	*	*	*	*	*	–
A*n* = –A*n*, A*n* =– A*n* ;	*	*	*	*	*	*	–
A*n* = A*n* (S) ;	*	*	–	*	*	–	–
A*n* = A*n* (S), A*n* = A*n* (S) ;	*	*	–	*	*	–	–
Dreg_lo_hi = Dreg (RND) ;	*	*	–	–	–	*	–
Dreg_lo_hi = Dreg + Dreg (RND12) ;	*	*	–	–	–	*	–
Dreg_lo_hi = Dreg – Dreg (RND12) ;	*	*	–	–	–	*	–

Table 2-9. ALU Instruction Summary (Cont'd)

Instruction	ASTAT Status Flags						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	V V_COPY VS	AQ
Dreg_lo_hi = Dreg + Dreg (RND20) ;	*	*	–	–	–	*	–
Dreg_lo_hi = Dreg – Dreg (RND20) ;	*	*	–	–	–	*	–
Dreg_lo = SIGNBITS Dreg ;	–	–	–	–	–	–	–
Dreg_lo = SIGNBITS Dreg_lo_hi ;	–	–	–	–	–	–	–
Dreg_lo = SIGNBITS An ;	–	–	–	–	–	–	–
Dreg_lo = EXPADJ (Dreg, Dreg_lo) (V) ;	–	–	–	–	–	–	–
Dreg_lo = EXPADJ (Dreg_lo_hi, Dreg_lo);	–	–	–	–	–	–	–
Dreg = Dreg & Dreg ;	*	*	**	–	–	**/–	–
Dreg = ~ Dreg ;	*	*	**	–	–	**/–	–
Dreg = Dreg | Dreg ;	*	*	**	–	–	**/–	–
Dreg = Dreg ^ Dreg ;	*	*	**	–	–	**/–	–
Dreg =– Dreg ;	*	*	*	–	–	*	–

ALU Data Flow Details

Figure 2-9 shows a more detailed diagram of the Arithmetic Units and Data Register File, which appears in Figure 2-1 on page 2-2.

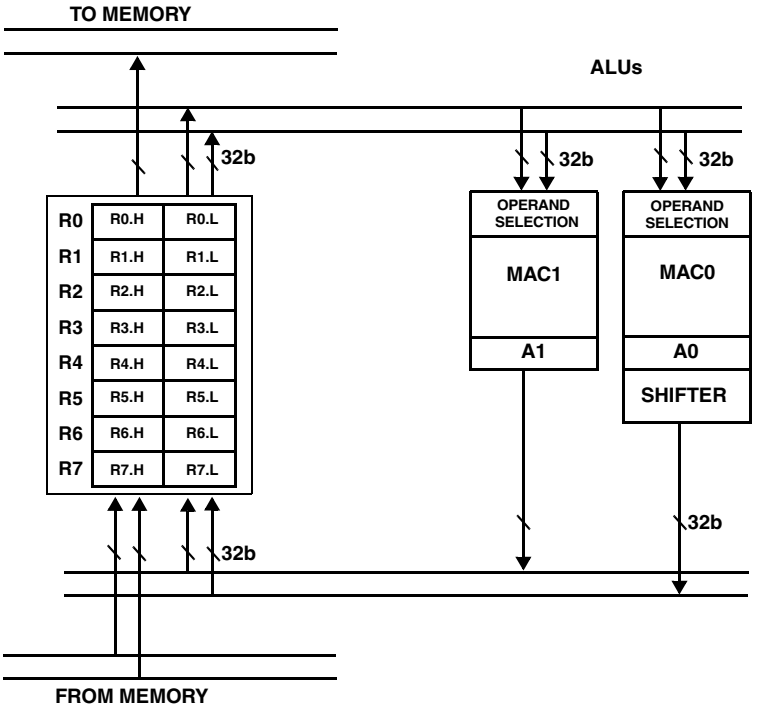

Figure 2-9. Register Files and ALUs

ALU0 is described here for convenience. ALU1 is very similar—a subset of ALU0.

Each ALU performs 40-bit addition for the accumulation of the multiplier results, as well as 32-bit and dual 16-bit operations. Each ALU has two 32-bit input ports that can be considered a pair of 16-bit operands or a single 32-bit operand. For single 16-bit operations, any of the four

possible 16-bit operands may be used with any of the other 16-bit operands presented at the input to the ALU.

As shown in Figure 2-10, for dual 16-bit operations, the high halves and low halves are paired, providing four possible combinations of addition and subtraction.

```
(A) H + H, L + L    (B) H + H, L - L
(C) H - H, L + L    (D) H - H, L - L
```

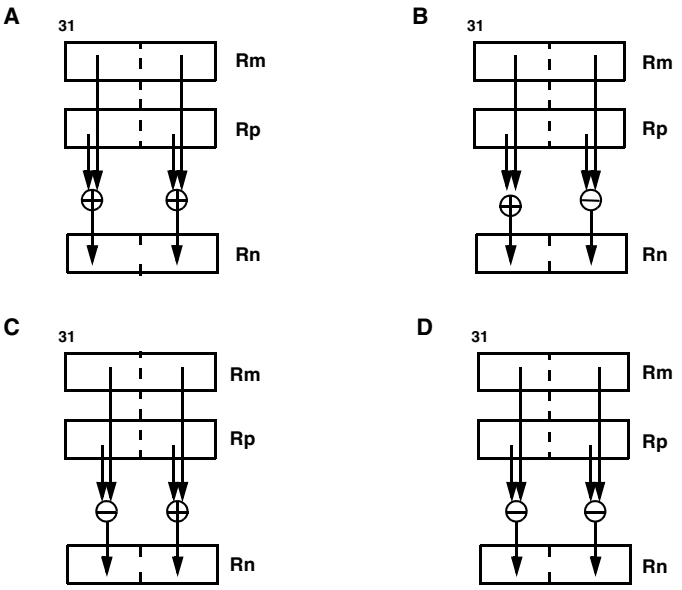

Figure 2-10. Dual 16-Bit ALU Operations

Dual 16-Bit Cross Options

For dual 16-bit operations, the results may be crossed. “Crossing the results” changes the location in the result register for the result of a calculation. Usually, the result from the high side calculation is placed in the

high half of the result register, and the result from the low side calculation is placed in the low half of the result register. With the cross option, the high result is placed in the low half of the destination register, and the low result is placed in the high half of the destination register (see Figure 2-11). This is particularly useful when dealing with complex math and portions of the Fast Fourier Transform (FFT). The cross option applies to ALU0 only.

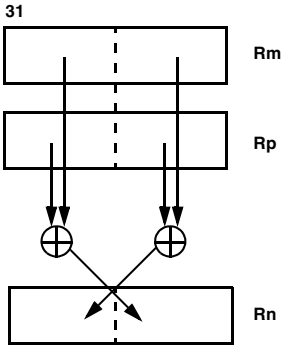

Figure 2-11. Cross Options for Dual 16-Bit ALU Operations

ALU Status Signals

Each ALU generates six status signals: the zero (AZ) status, the negative (AN) status, the carry (ACn) status, the sticky overflow (AVnS) status, the immediate overflow (AVn) status, and the quotient (AQ) status. All arithmetic status signals are latched into the arithmetic status register (ASTAT) at the end of the cycle. For the effect of ALU instructions on the status flags, see Table 2-9 on page 2-30.

Depending on the instruction, the inputs can come from the Data Register File, the Pointer Register File, or the Arithmetic Result registers. Arithmetic on 32-bit operands directly support multi precision operations in the ALU.

ALU Division Support Features

The ALU supports division with two special divide primitives. These instructions (`DIVS`, `DIVQ`) let programs implement a non-restoring, conditional (error checking), add/subtract/division algorithm.

The division can be either signed or unsigned, but both the dividend and divisor must be of the same type. Details about using division and programming examples are available in the *Blackfin Processor Programming Reference*.

Special SIMD Video ALU Operations

Four 8-bit Video ALUs enable the processor to process video information with high efficiency. Each Video ALU instruction may take from one to four pairs of 8-bit inputs and return one to four 8-bit results. The inputs are presented to the Video ALUs in two 32-bit words from the Data Register File. The possible operations include:

- Quad 8-bit add or subtract
- Quad 8-bit average
- Quad 8-bit pack or unpack
- Quad 8-bit subtract-absolute-accumulate
- Byte align

For more information about the operation of these instructions, see the *Blackfin Processor Programming Reference*.

Multiply Accumulators (Multipliers)

The two multipliers (MAC0 and MAC1) perform fixed-point multiplication and multiply and accumulate operations. Multiply and accumulate

operations are available with either cumulative addition or cumulative subtraction.

Multiplier fixed-point instructions operate on 16-bit fixed-point data and produce 32-bit results that may be added or subtracted from a 40-bit accumulator.

Inputs are treated as fractional or integer, unsigned or two's-complement. Multiplier instructions include:

- Multiplication
- Multiply and accumulate with addition, rounding optional
- Multiply and accumulate with subtraction, rounding optional
- Dual versions of the above

Multiplier Operation

Each multiplier has two 32-bit inputs from which it derives the two 16-bit operands. For single multiply and accumulate instructions, these operands can be any Data registers in the Data Register File. Each multiplier can accumulate results in its Accumulator register, `A1` or `A0`. The accumulator results can be saturated to 32 or 40 bits. The multiplier result can also be written directly to a 16- or 32-bit destination register with optional rounding.

Each multiplier instruction determines whether the inputs are either both in integer format or both in fractional format. The format of the result matches the format of the inputs. In MAC0, both inputs are treated as signed or unsigned. In MAC1, there is a mixed-mode option.

If both inputs are fractional and signed, the multiplier automatically shifts the result left one bit to remove the redundant sign bit. Unsigned fractional, integer, and mixed modes do not perform a shift for sign bit correction. Multiplier instruction options specify the data format of the

inputs. See "Multiplier Instruction Options" on page 2-41 for more information.

Placing Multiplier Results in Multiplier Accumulator Registers

As shown in Figure 2-9 on page 2-33, each multiplier has a dedicated accumulator, `A0` or `A1`. Each Accumulator register is divided into three sections—`A0.L/A1.L` (bits 15:0), `A0.H/A1.H` (bits 31:16), and `A0.X/A1.X` (bits 39:32).

When the multiplier writes to its result Accumulator registers, the 32-bit result is deposited into the lower bits of the combined Accumulator register, and the MSB is sign-extended into the upper eight bits of the register (`A0.X/A1.X`).

Multiplier output can be deposited not only in the `A0` or `A1` registers, but also in a variety of 16- or 32-bit Data registers in the Data Register File.

Rounding or Saturating Multiplier Results

On a multiply and accumulate operation, the accumulator data can be saturated and, optionally, rounded for extraction to a register or register half. When a multiply deposits a result only in a register or register half, the saturation and rounding works the same way. The rounding and saturation operations work as follows.

- Rounding is applied only to fractional results except for the `IH` option, which applies rounding and high half extraction to an integer result.

  For the IH option, the rounded result is obtained by adding 0x8000 to the accumulator (for MAC) or multiply result (for mult) and then saturating to 32-bits. For more information, see "Rounding Multiplier Results" on page 2-17.

- If an overflow or underflow has occurred, the saturate operation sets the specified Result register to the maximum positive or negative value. For more information, see the following section.

Saturating Multiplier Results on Overflow

The following bits in ASTAT indicate multiplier overflow status:

- Bit 16 (AV0) and bit 18 (AV1) record overflow condition (whether the result has overflowed 32 bits) for the A0 and A1 accumulators, respectively.
- If the bit is cleared (=0), no overflow or underflow has occurred. If the bit is set (=1), an overflow or underflow has occurred. The AV0S and AV1S bits are sticky bits.
- Bit 24 (V) and bit 25 (VS) are set if overflow occurs in extracting the accumulator result to a register.

Multiplier Instruction Summary

Table 2-10 lists the multiplier instructions. For more information about assembly language syntax and the effect of multiplier instructions on the status flags, see the *Blackfin Processor Programming Reference*.

In Table 2-10, note the meaning of these symbols:

- Dreg denotes any Data Register File register.
- Dreg_lo_hi denotes any 16-bit register half in any Data Register File register.
- Dreg_lo denotes the lower 16 bits of any Data Register File register.
- Dreg_hi denotes the upper 16 bits of any Data Register File register.

- A*n* denotes either MAC Accumulator register `A0` or `A1`.
- * Indicates the flag may be set or cleared, depending on the results of the instruction.
- – Indicates no effect.

Table 2-10. Multiplier Instruction Summary

Instruction	ASTAT Status Flags		
	AV0 AV0S	AV1 AV1S	V V_COPY VS
Dreg_lo = Dreg_lo_hi * Dreg_lo_hi ;	–	–	*
Dreg_hi = Dreg_lo_hi * Dreg_lo_hi ;	–	–	*
Dreg = Dreg_lo_hi * Dreg_lo_hi ;	–	–	*
A*n* = Dreg_lo_hi * Dreg_lo_hi ;	*	*	–
A*n* += Dreg_lo_hi * Dreg_lo_hi ;	*	*	–
An –= Dreg_lo_hi * Dreg_lo_hi ;	*	*	–
Dreg_lo = (A0 = Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg_lo = (A0 += Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg_lo = (A0 –= Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg_hi = (A1 = Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg_hi = (A1 += Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg_hi = (A1 –= Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg = (A*n* = Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg = (A*n* += Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg = (A*n* –= Dreg_lo_hi * Dreg_lo_hi) ;	*	*	*
Dreg *= Dreg ;	–	–	–

Multiplier Instruction Options

The descriptions of multiplier instruction options are shown in Table 2-11. Not all options are available for all instructions. For information about how to use these options with their respective instructions, see the *Blackfin Processor Programming Reference*.

Table 2-11. Multiplier Instruction Options

default	No option; input data is signed fraction.
(IS)	Input data operands are signed integer. No shift correction is made.
(FU)	Input data operands are unsigned fraction. No shift correction is made.
(IU)	Input data operands are unsigned integer. No shift correction is made.
(T)	Input data operands are signed fraction. When copying to the destination half register, truncates the lower 16 bits of the Accumulator contents.
(TFU)	Input data operands are unsigned fraction. When copying to the destination half register, truncates the lower 16 bits of the Accumulator contents.
(S2RND)	If multiplying and accumulating to a register: Input data operands are signed fraction. When copying to the destination register, Accumulator contents are scaled (multiply x2 by a one-place shift-left) and rounded. If scaling and rounding produce a signed value larger than 32 bits, the number is saturated to its maximum positive or negative value. If multiplying and accumulating to a register: Input data operands are signed fraction. When copying to the destination register, Accumulator contents are scaled (multiplied x2 by a one-place shift-left) and the upper 16 bits are rounded before truncating the lower 16 bits of the Accumulator. If scaling and rounding produce a signed value larger than 16 bits, the number is saturated to its maximum positive or negative value.
(ISS2)	If multiplying and accumulating to a register: Input data operands are signed integer. When copying to the destination register, Accumulator contents are scaled (multiplied x2 by a one-place shift-left). If scaling produces a signed value larger than 32 bits, the number is saturated to its maximum positive or negative value. If multiplying and accumulating to a half register: When copying the lower 16 bits to the destination half register, the Accumulator contents are scaled. If scaling produces a signed value greater than 16 bits, the number is saturated to its maximum positive or negative value.

Table 2-11. Multiplier Instruction Options (Cont'd)

(IH)	This option indicates integer multiplication with high half word extraction. The Accumulator is saturated at 32 bits, and bits [31:16] of the Accumulator are rounded, and then copied into the destination half register.
(W32)	Input data operands are signed fraction with no extension bits in the Accumulators at 32 bits. Left-shift correction of the product is performed, as required. This option is used for legacy GSM speech vocoder algorithms written for 32-bit Accumulators. For this option only, this special case applies: `0x8000` × `0x8000 = 0x7FFF`.
(M)	Operation uses mixed-multiply mode. Valid only for MAC1 versions of the instruction. Multiplies a signed fraction by an unsigned fractional operand with no left-shift correction. Operand one is signed; operand two is unsigned. MAC0 performs an unmixed multiply on signed fractions by default, or another format as specified. That is, MAC0 executes the specified signed/signed or unsigned/unsigned multiplication. The `(M)` option can be used alone or in conjunction with one other format option.

Multiplier Data Flow Details

Figure 2-12 shows the Register files and ALUs, along with the multiplier/accumulators.

Each multiplier has two 16-bit inputs, performs a 16-bit multiplication, and stores the result in a 40-bit accumulator or extracts to a 16-bit or 32-bit register. Two 32-bit words are available at the MAC inputs, providing four 16-bit operands to chose from.

One of the operands must be selected from the low half or the high half of one 32-bit word. The other operand must be selected from the low half or the high half of the other 32-bit word. Thus, each MAC is presented with four possible input operand combinations. The two 32-bit words can contain the same register information, giving the options for squaring and multiplying the high half and low half of the same register. Figure 2-13 show these possible combinations.

The 32-bit product is passed to a 40-bit adder/subtracter, which may add or subtract the new product from the contents of the Accumulator Result register or pass the new product directly to the Data Register File Results

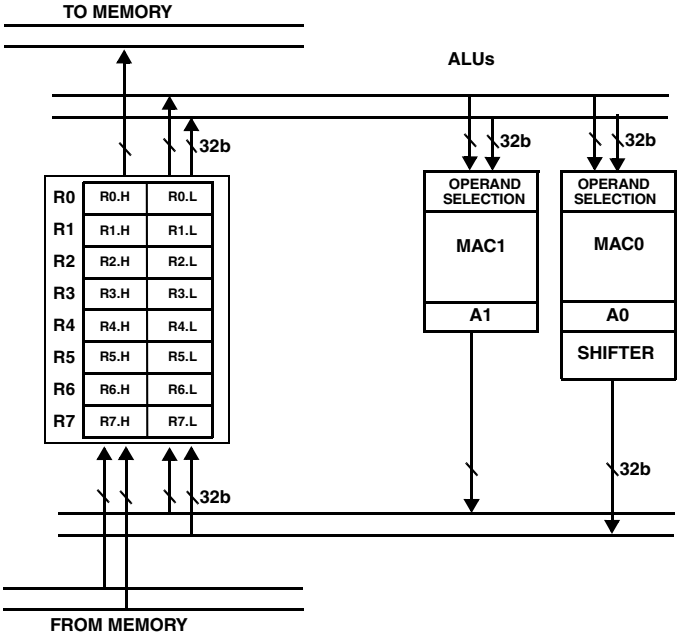

Figure 2-12. Register Files and ALUs

register. For results, the A0 and A1 registers are 40 bits wide. Each of these registers consists of smaller 32- and 8-bit registers—A0.W, A1.W, A0.X, and A1.X.

Some example instructions:

```
A0 = R3.L * R4.H ;
```

In this instruction, the MAC0 multiplier/accumulator performs a multiply and puts the result in the Accumulator register.

```
A1 += R3.H * R4.H ;
```

In this instruction, the MAC1 multiplier/accumulator performs a multiply and accumulates the result with the previous results in the A1 Accumulator.

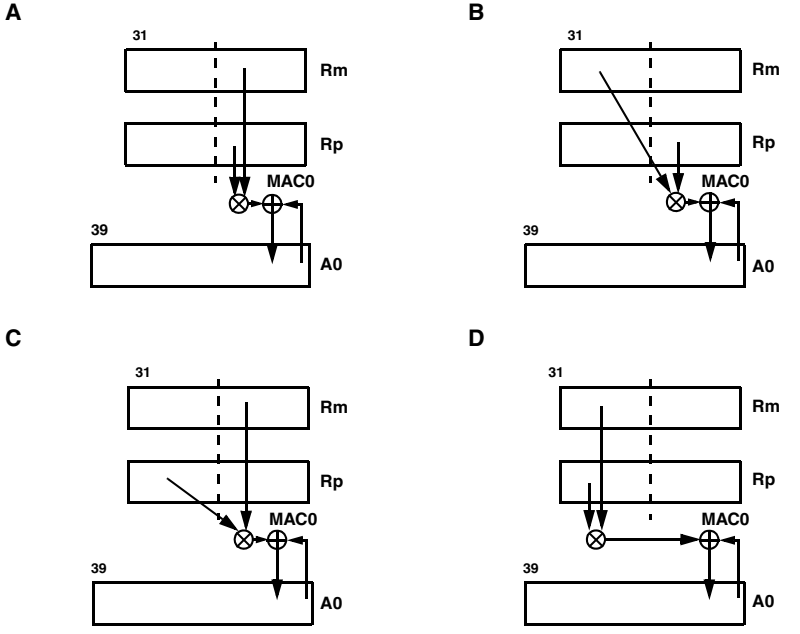

Figure 2-13. Four Possible Combinations of MAC Operations

Multiply Without Accumulate

The multiplier may operate without the accumulation function. If accumulation is not used, the result can be directly stored in a register from the Data Register File or the Accumulator register. The destination register may be 16 bits or 32 bits. If a 16-bit destination register is a low half, then MAC0 is used; if it is a high half, then MAC1 is used. For a 32-bit destination register, either MAC0 or MAC1 is used.

If the destination register is 16 bits, then the word that is extracted from the multiplier depends on the data type of the input.

- If the multiplication uses fractional operands or the `IH` option, then the high half of the result is extracted and stored in the 16-bit destination registers (see Figure 2-14).
- If the multiplication uses integer operands, then the low half of the result is extracted and stored in the 16-bit destination registers. These extractions provide the most useful information in the resultant 16-bit word for the data type chosen (see Figure 2-15).

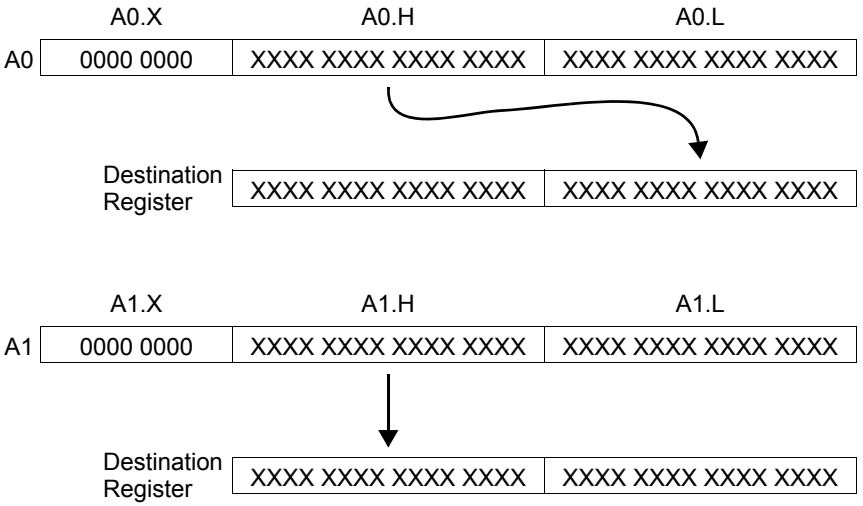

Figure 2-14. Multiplication of Fractional Operands

For example, this instruction uses fractional, unsigned operands:

```
R0.L = R1.L * R2.L (FU) ;
```

The instruction deposits the upper 16 bits of the multiply answer with rounding and saturation into the lower half of R0, using MAC0. This instruction uses unsigned integer operands:

```
R0.H = R2.H * R3.H (IU) ;
```

The instruction deposits the lower 16 bits of the multiply answer with any required saturation into the high half of R0, using MAC1.

```
R0 = R1.L * R2.L ;
```

Regardless of operand type, the preceding operation deposits 32 bits of the multiplier answer with saturation into R0, using MAC0.

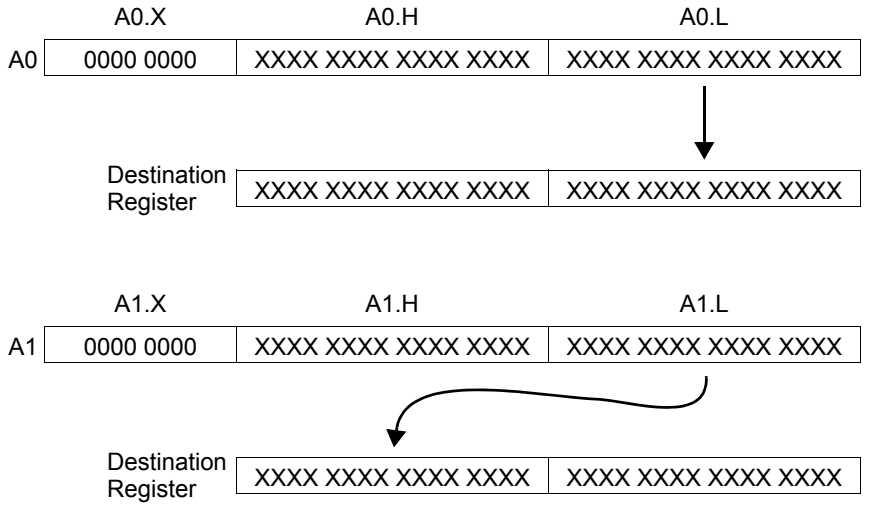

Figure 2-15. Multiplication of Integer Operands

Special 32-Bit Integer MAC Instruction

The processor supports a multicycle 32-bit MAC instruction:

```
Dreg *= Dreg
```

The single instruction multiplies two 32-bit integer operands and provides a 32-bit integer result, destroying one of the input operands.

The instruction takes multiple cycles to execute. Refer to the product data sheet and the *Blackfin Processor Programming Reference* for more information about the exact operation of this instruction. This macro function is interruptable and does not modify the data in either Accumulator register `A0` or `A1`.

Dual MAC Operations

The processor has two 16-bit MACs. Both MACs can be used in the same operation to double the MAC throughput. The same two 32-bit input registers are offered to each MAC unit, providing each with four possible combinations of 16-bit input operands. Dual MAC operations are frequently referred to as vector operations, because a program could store vectors of samples in the four input operands and perform vector computations.

An example of a dual multiply and accumulate instruction is

```
A1 += R1.H * R2.L, A0 += R1.L * R2.H ;
```

This instruction represents two multiply and accumulate operations.

- In one operation (MAC1) the high half of `R1` is multiplied by the low half of `R2` and added to the contents of the `A1` Accumulator.
- In the second operation (MAC0) the low half of `R1` is multiplied by the high half of `R2` and added to the contents of `A0`.

The results of the MAC operations may be written to registers in a number of ways: as a pair of 16-bit halves, as a pair of 32-bit registers, or as an independent 16-bit half register or 32-bit register.

For example:

```
R3.H = (A1 += R1.H * R2.L), R3.L = (A0 += R1.L * R2.L) ;
```

In this instruction, the 40-bit Accumulator is packed into a 16-bit half register. The result from MAC1 must be transferred to a high half of a destination register and the result from MAC0 must be transferred to the low half of the same destination register.

The operand type determines the correct bits to extract from the Accumulator and deposit in the 16-bit destination register. See "Multiply Without Accumulate" on page 2-44.

```
R3 = (A1 += R1.H * R2.L), R2 = (A0 += R1.L * R2.L) ;
```

In this instruction, the 40-bit Accumulators are packed into two 32-bit registers. The registers must be register pairs (`R[1:0]`, `R[3:2]`, `R[5:4]`, `R[7:6]`).

```
R3.H = (A1 += R1.H * R2.L), A0 += R1.L * R2.L ;
```

This instruction is an example of one Accumulator—but not the other—being transferred to a register. Either a 16- or 32-bit register may be specified as the destination register.

Barrel Shifter (Shifter)

The shifter provides bitwise shifting functions for 16-, 32-, or 40-bit inputs, yielding a 16-, 32-, or 40-bit output. These functions include arithmetic shift, logical shift, rotate, and various bit test, set, pack, unpack, and exponent detection functions. These shift functions can be combined to implement numerical format control, including full floating-point representation.

Shifter Operations

The shifter instructions (`>>>`, `>>`, `<<`, `ASHIFT`, `LSHIFT`, `ROT`) can be used various ways, depending on the underlying arithmetic requirements. The `ASHIFT` and `>>>` instructions represent the arithmetic shift. The `LSHIFT`, `<<`, and `>>` instructions represent the logical shift.

The arithmetic shift and logical shift operations can be further broken into subsections. Instructions that are intended to operate on 16-bit single or paired numeric values (as would occur in many DSP algorithms) can use the instructions `ASHIFT` and `LSHIFT`. These are typically three-operand instructions.

Instructions that are intended to operate on a 32-bit register value and use two operands, such as instructions frequently used by a compiler, can use the `>>>` and `>>` instructions.

Arithmetic shift, logical shift, and rotate instructions can obtain the shift argument from a register or directly from an immediate value in the instruction. For details about shifter related instructions, see "Shifter Instruction Summary" on page 2-52.

Two-Operand Shifts

Two-operand shift instructions shift an input register and deposit the result in the same register.

Immediate Shifts

An immediate shift instruction shifts the input bit pattern to the right (downshift) or left (upshift) by a given number of bits. Immediate shift instructions use the data value in the instruction itself to control the amount and direction of the shifting operation.

The following example shows the input value downshifted.

```
R0 contains 0000 B6A3 ;
R0 >>= 0x04 ;
```

results in

```
R0 contains 0000 0B6A ;
```

The following example shows the input value upshifted.

```
R0 contains 0000 B6A3 ;
R0 <<= 0x04 ;
```

results in

```
R0 contains 000B 6A30 ;
```

Register Shifts

Register-based shifts use a register to hold the shift value. The entire 32-bit register is used to derive the shift value, and when the magnitude of the shift is greater than or equal to 32, then the result is either 0 or –1.

The following example shows the input value upshifted.

```
R0 contains 0000 B6A3 ;
R2 contains 0000 0004 ;
R0 <<= R2 ;
```

results in

```
R0 contains 000B 6A30 ;
```

Three-Operand Shifts

Three-operand shifter instructions shift an input register and deposit the result in a destination register.

Immediate Shifts

Immediate shift instructions use the data value in the instruction itself to control the amount and direction of the shifting operation.

The following example shows the input value downshifted.

```
R0 contains 0000 B6A3 ;
R1 = R0 >> 0x04 ;
```

results in

```
R1 contains 0000 0B6A ;
```

The following example shows the input value upshifted.

```
R0.L contains B6A3 ;
R1.H = R0.L << 0x04 ;
```

results in

```
R1.H contains 6A30 ;
```

Register Shifts

Register-based shifts use a register to hold the shift value. When a register is used to hold the shift value (for `ASHIFT`, `LSHIFT` or `ROT`), then the shift value is always found in the low half of a register (`Rn.L`). The bottom six bits of `Rn.L` are masked off and used as the shift value.

The following example shows the input value upshifted.

```
R0 contains 0000 B6A3 ;
R2.L contains 0004 ;
R1 = R0 ASHIFT by R2.L ;
```

results in

```
R1 contains 000B 6A30 ;
```

The following example shows the input value rotated. Assume the Condition Code (`CC`) bit is set to 0. For more information about `CC`, see "Condition Code Flag" on page 4-12.

```
R0 contains ABCD EF12 ;
R2.L contains 0004 ;
R1 = R0 ROT by R2.L ;
```

results in

```
R1 contains BCDE F125 ;
```

Note the `CC` bit is included in the result, at bit 3.

Bit Test, Set, Clear, Toggle

The shifter provides the method to test, set, clear, and toggle specific bits of a data register. All instructions have two arguments—the source register and the bit field value. The test instruction does not change the source register. The result of the test instruction resides in the CC bit.

The following examples show a variety of operations.

```
BITCLR (R0, 6) ;
BITSET (R2, 9) ;
BITTGL (R3, 2) ;
CC = BITTST (R3, 0) ;
```

Field Extract and Field Deposit

If the shifter is used, a source field may be deposited anywhere in a 32-bit destination field. The source field may be from 1 bit to 16 bits in length. In addition, a 1- to 16-bit field may be extracted from anywhere within a 32-bit source field.

Two register arguments are used for these functions. One holds the 32-bit destination or 32-bit source. The other holds the extract/deposit value, its length, and its position within the source.

Shifter Instruction Summary

Table 2-12 lists the shifter instructions. For more information about assembly language syntax and the effect of shifter instructions on the status flags, see the *Blackfin Processor Programming Reference*.

In Table 2-12, note the meaning of these symbols:

- Dreg denotes any Data Register File register.
- Dreg_lo denotes the lower 16 bits of any Data Register File register.

- Dreg_hi denotes the upper 16 bits of any Data Register File register.
- * Indicates the flag may be set or cleared, depending on the results of the instruction.
- * 0 Indicates versions of the instruction that send results to Accumulator `A0` set or clear `AV0`.
- * 1 Indicates versions of the instruction that send results to Accumulator `A1` set or clear `AV1`.
- ** Indicates the flag is cleared.
- *** Indicates `CC` contains the latest value shifted into it.
- – Indicates no effect.

Table 2-12. Shifter Instruction Summary

Instruction	ASTAT Status Flag						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	CC	V V_COPY VS
BITCLR (Dreg, uimm5) ;	*	*	**	–	–	–	**/–
BITSET (Dreg, uimm5) ;	**	*	**	–	–	–	**/–
BITTGL (Dreg, uimm5) ;	*	*	**	–	–	–	**/–
CC = BITTST (Dreg, uimm5) ;	–	–	–	–	–	*	–
CC = !BITTST (Dreg, uimm5) ;	–	–	–	–	–	*	–
Dreg = DEPOSIT (Dreg, Dreg) ;	*	*	**	–	–	–	**/–
Dreg = EXTRACT (Dreg, Dreg) ;	*	*	**	–	–	–	**/–
BITMUX (Dreg, Dreg, A0) ;	–	–	–	–	–	–	–

Table 2-12. Shifter Instruction Summary (Cont'd)

Instruction	ASTAT Status Flag						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	CC	V V_COPY VS
Dreg_lo = ONES Dreg ;	–	–	–	–	–	–	–
Dreg = PACK (Dreg_lo_hi, Dreg_lo_hi);	–	–	–	–	–	–	–
Dreg >>>= uimm5 ;	*	*	–	–	–	–	**/–
Dreg >>= uimm5 ;	*	*	–	–	–	–	**/–
Dreg <<= uimm5 ;	*	*	–	–	–	–	**/–
Dreg = Dreg >>> uimm5 ;	*	*	–	–	–	–	**/–
Dreg = Dreg >> uimm5 ;	*	*	–	–	–	–	**/–
Dreg = Dreg << uimm5 ;	*	*	–	–	–	–	*
Dreg = Dreg >>> uimm4 (V) ;	*	*	–	–	–	–	**/–
Dreg = Dreg >> uimm4 (V) ;	*	*	–	–	–	–	**/–
Dreg = Dreg << uimm4 (V) ;	*	*	–	–	–	–	*
A*n* = A*n* >>> uimm5 ;	*	*	–	** 0/–	** 1/–	–	–
A*n* = A*n* >> uimm5 ;	*	*	–	** 0/–	** 1/–	–	–
A*n* = A*n* << uimm5 ;	*	*	–	* 0	* 1	–	–
Dreg_lo_hi = Dreg_lo_hi >>> uimm4 ;	*	*	–	–	–	–	**/–
Dreg_lo_hi = Dreg_lo_hi >> uimm4 ;	*	*	–	–	–	–	**/–
Dreg_lo_hi = Dreg_lo_hi << uimm4 ;	*	*	–	–	–	–	*
Dreg >>>= Dreg ;	*	*	–	–	–	–	**/–
Dreg >>= Dreg ;	*	*	–	–	–	–	**/–
Dreg <<= Dreg ;	*	*	–	–	–	–	**/–

Table 2-12. Shifter Instruction Summary (Cont'd)

Instruction	ASTAT Status Flag						
	AZ	AN	AC0 AC0_COPY AC1	AV0 AV0S	AV1 AV1S	CC	V V_COPY VS
Dreg = ASHIFT Dreg BY Dreg_lo ;	*	*	–	–	–	–	*
Dreg = LSHIFT Dreg BY Dreg_lo ;	*	*	–	–	–	–	**/–
Dreg = ROT Dreg BY imm6 ;	–	–	–	–	–	***	–
Dreg = ASHIFT Dreg BY Dreg_lo (V) ;	*	*	–	–	–	–	*
Dreg = LSHIFT Dreg BY Dreg_lo (V) ;	*	*	–	–	–	–	**/–
Dreg_lo_hi = ASHIFT Dreg_lo_hi BY Dreg_lo ;	*	*	–	–	–	–	*
Dreg_lo_hi = LSHIFT Dreg_lo_hi BY Dreg_lo ;	*	*	–	–	–	–	**/–
A*n* = A*n* ASHIFT BY Dreg _lo ;	*	*	–	* 0	* 1	–	–
A*n* = A*n* ROT BY imm6 ;	–	–	–	–	–	***	–
Preg = Preg >> 1 ;	–	–	–	–	–	–	–
Preg = Preg >> 2 ;	–	–	–	–	–	–	–
Preg = Preg << 1 ;	–	–	–	–	–	–	–
Preg = Preg << 2 ;	–	–	–	–	–	–	–
Dreg = (Dreg + Dreg) << 1 ;	*	*	*	–	–	–	*
Dreg = (Dreg + Dreg) << 2 ;	*	*	*	–	–	–	*
Preg = (Preg + Preg) << 1 ;	–	–	–	–	–	–	–
Preg = (Preg + Preg) << 2 ;	–	–	–	–	–	–	–
Preg = Preg + (Preg << 1) ;	–	–	–	–	–	–	–
Preg = Preg + (Preg << 2) ;	–	–	–	–	–	–	–

3 OPERATING MODES AND STATES

Blackfin processors support the following three processor modes:

- User mode
- Supervisor mode
- Emulation mode

Emulation and Supervisor modes have unrestricted access to the core resources. User mode has restricted access to certain system resources, thus providing a protected software environment.

User mode is considered the domain of application programs. Supervisor mode and Emulation mode are usually reserved for the kernel code of an operating system.

The processor mode is determined by the Event Controller. When servicing an interrupt, a nonmaskable interrupt (NMI), or an exception, the processor is in Supervisor mode. When servicing an emulation event, the processor is in Emulation mode. When not servicing any events, the processor is in User mode.

The current processor mode may be identified by interrogating the IPEND memory-mapped register (MMR), as shown in Table 3-1.

MMRs cannot be read while the processor is in User mode.

Table 3-1. Identifying the Current Processor Mode

Event	Mode	IPEND
Interrupt	Supervisor	≥ 0x10 but IPEND[0], IPEND[1], IPEND[2], and IPEND[3] = 0.
Exception	Supervisor	≥ 0x08 The core is processing an exception event if IPEND[0] = 0, IPEND[1] = 0, IPEND[2] = 0, IPEND[3] = 1, and IPEND[15:4] are 0's or 1's.
NMI	Supervisor	≥ 0x04 The core is processing an NMI event if IPEND[0] = 0, IPEND[1] = 0, IPEND[2] = 1, and IPEND[15:2] are 0's or 1's.
Reset	Supervisor	= 0x02 As the reset state is exited, IPEND is set to 0x02, and the reset vector runs in Supervisor mode.
Emulation	Emulator	= 0x01 The processor is in Emulation mode if IPEND[0] = 1, regardless of the state of the remaining bits IPEND[15:1].
None	User	= 0x00

In addition, the processor supports the following two non-processing states:

- Idle state
- Reset state

Figure 3-1 illustrates the processor modes and states as well as the transition conditions between them.

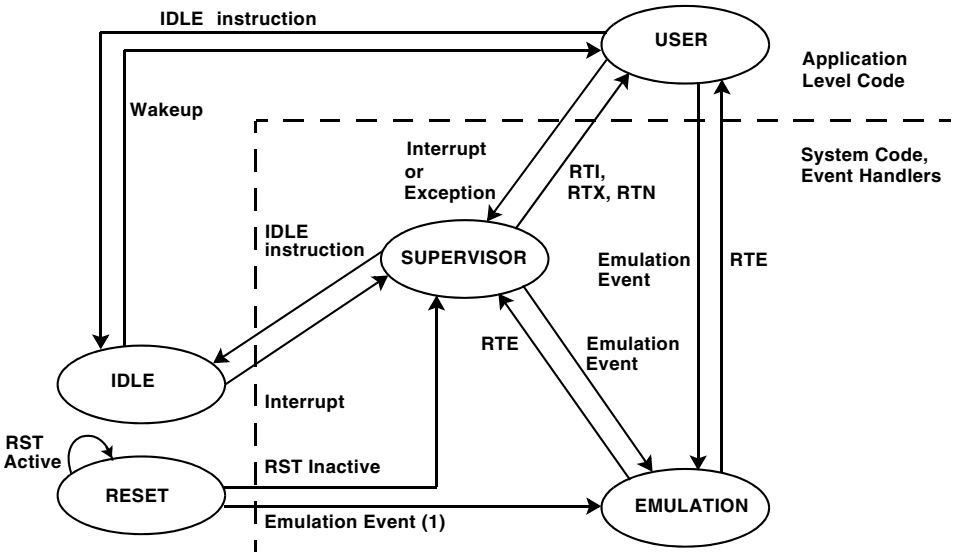

(1) Normal exit from Reset is to Supervisor mode. However, emulation hardware may have initiated a reset. If so, exit from Reset is to Emulation.

Figure 3-1. Processor Modes and States

User Mode

The processor is in User mode when it is not in Reset or Idle state, and when it is not servicing an interrupt, NMI, exception, or emulation event. User mode is used to process application level code that does not require explicit access to system registers. Any attempt to access restricted system registers causes an exception event. Table 3-2 lists the registers that may be accessed in User mode.

Table 3-2. Registers Accessible in User Mode

Processor Registers	Register Names
Data Registers	R[7:0], A[1:0]
Pointer Registers	P[5:0], SP, FP, I[3:0], M[3:0], L[3:0], B[3:0]
Sequencer and Status Registers	RETS, LC[1:0], LT[1:0], LB[1:0], ASTAT, CYCLES, CYCLES2

Protected Resources and Instructions

System resources consist of a subset of processor registers, all MMRs, and a subset of protected instructions. These system and core MMRs are located starting at address 0xFFC0 0000. This region of memory is protected from User mode access. Any attempt to access MMR space in User mode causes an exception.

A list of protected instructions appears in Table 3-3. Any attempt to issue any of the protected instructions from User mode causes an exception event.

Table 3-3. Protected Instructions

Instruction	Description
RTI	Return from Interrupt
RTX	Return from Exception
RTN	Return from NMI
CLI	Disable Interrupts
STI	Enable Interrupts
RAISE	Force Interrupt/Reset
RTE	Return from Emulation Causes an exception only if executed outside Emulation mode

Protected Memory

Additional memory locations can be protected from User mode access. A Cacheability Protection Lookaside Buffer (CPLB) entry can be created and enabled. See "Memory Management Unit" on page 6-47 for further information.

Entering User Mode

When coming out of reset, the processor is in Supervisor mode because it is servicing a reset event. To enter User mode from the Reset state, two steps must be performed. First, a return address must be loaded into the `RETI` register. Second, an `RTI` must be issued. The following example code shows how to enter User mode upon reset.

Example Code to Enter User Mode Upon Reset

Listing 3-1 provides code for entering User mode from reset.

Listing 3-1. Entering User Mode from Reset

```
P1.L = START ;   /* Point to start of user code */
P1.H = START ;
RETI = P1 ;
RTI ;   /* Return from Reset Event */

START :   /* Place user code here */
```

Return Instructions That Invoke User Mode

Table 3-4 provides a summary of return instructions that can be used to invoke User mode from various processor event service routines. When these instructions are used in service routines, the value of the return address must be first stored in the appropriate event `RETx` register. In the

case of an interrupt routine, if the service routine is interruptible, the return address is stored on the stack. For this case, the address can be found by popping the value from the stack into RETI. Once RETI has been loaded, the RTI instruction can be issued.

Note the stack pop is optional. If the RETI register is not pushed/popped, then the interrupt service routine becomes non-interruptible, because the return address is not saved on the stack.

The processor remains in User mode until one of these events occurs:

- An interrupt, NMI, or exception event invokes Supervisor mode.
- An emulation event invokes Emulation mode.
- A reset event invokes the Reset state.

Table 3-4. Return Instructions That Can Invoke User Mode

Current Process Activity	Return Instruction to Use	Execution Resumes at Address in This Register
Interrupt Service Routine	RTI	RETI
Exception Service Routine	RTX	RETX
Nonmaskable Interrupt Service Routine	RTN	RETN
Emulation Service Routine	RTE	RETE

Supervisor Mode

The processor services all interrupt, NMI, and exception events in Supervisor mode.

Supervisor mode has full, unrestricted access to all processor system resources, including all emulation resources, unless a CPLB has been

configured and enabled. See "Memory Management Unit" on page 6-47 for a further description. Only Supervisor mode can use the register alias USP, which references the User Stack Pointer in memory. This register alias is necessary because in Supervisor mode, SP refers to the kernel stack pointer rather than to the user stack pointer.

Normal processing begins in Supervisor mode from the Reset state. De-asserting the RESET signal switches the processor from the Reset state to Supervisor mode where it remains until an emulation event or Return instruction occurs to change the mode. Before the Return instruction is issued, the RETI register must be loaded with a valid return address.

Non-OS Environments

For non-OS environments, application code should remain in Supervisor mode so that it can access all core and system resources. When $\overline{\text{RESET}}$ is de-asserted, the processor initiates operation by servicing the reset event. Emulation is the only event that can pre-empt this activity. Therefore, lower priority events cannot be processed.

One way of keeping the processor in Supervisor mode and still allowing lower priority events to be processed is to set up and force the lowest priority interrupt (IVG15). Events and interrupts are described further in "Events and Sequencing" on page 4-18. After the low priority interrupt has been forced using the RAISE 15 instruction, RETI can be loaded with a return address that points to user code that can execute until IVG15 is issued. After RETI has been loaded, the RTI instruction can be issued to return from the reset event.

The interrupt handler for IVG15 can be set to jump to the application code starting address. An additional RTI is not required. As a result, the processor remains in Supervisor mode because IPEND[15] remains set. At this point, the processor is servicing the lowest priority interrupt. This ensures that higher priority interrupts can be processed.

Example Code to Stay in Supervisor Mode Coming Out of Reset

To remain in Supervisor mode when coming out of the Reset state, use code as shown in Listing 3-2.

Listing 3-2. Staying in Supervisor Mode Coming Out of Reset

```
P0.L = LO(EVT15) ;   /* Point to IVG15 in Event Vector Table */
P0.H = HI(EVT15) ;
P1.L = START ;   /* Point to start of User code */

P1.H = START ;
[P0] = P1 ;   /* Place the address of start code in IVG15 of EVT
*/

P0.L = LO(IMASK) ;

R0 = [P0] ;
R1.L = EVT_IVG15 & 0xFFFF ;

R0 = R0 | R1 ;
[P0] = R0 ;   /* Set (enable) IVG15 bit in Interrupt Mask Register
*/

RAISE 15 ;   /* Invoke IVG15 interrupt */
P0.L = WAIT_HERE ;
P0.H = WAIT_HERE ;
RETI = P0 ;   /* RETI loaded with return address */
RTI ;   /* Return from Reset Event */
WAIT_HERE :   /* Wait here till IVG15 interrupt is serviced */

JUMP WAIT_HERE ;
```

```
START:   /* IVG15 vectors here */
[--SP] = RETI ;   /* Enables interrupts and saves return address
to stack */
```

Emulation Mode

The processor enters Emulation mode if Emulation mode is enabled and either of these conditions is met:

- An external emulation event occurs.
- The `EMUEXCPT` instruction is issued.

The processor remains in Emulation mode until the emulation service routine executes an `RTE` instruction. If no interrupts are pending when the `RTE` instruction executes, the processor switches to User mode. Otherwise, the processor switches to Supervisor mode to service the interrupt.

Emulation mode is the highest priority mode, and the processor has unrestricted access to all system resources.

Idle State

Idle state stops all processor activity at the user's discretion, usually to conserve power during lulls in activity. No processing occurs during the Idle state. The Idle state is invoked by a sequential `IDLE` instruction. The `IDLE` instruction notifies the processor hardware that the Idle state is requested. An `SSYNC` instruction is performed automatically as part of the `IDLE` instruction to purge all speculative and transient states in the core and external system.

The processor remains in the Idle state until a peripheral or external device, such as a SPORT generates an interrupt that requires servicing.

In Listing 3-3, core interrupts are disabled and the IDLE instruction is executed. When all the pending processes have completed, the core disables its clocks. Since interrupts are disabled, Idle state can be terminated only by asserting a WAKEUP signal. For more information, see "System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26. (While not required, an interrupt could also be enabled in conjunction with the WAKEUP signal.)

When the WAKEUP signal is asserted, the processor wakes up, and the STI instruction enables interrupts again.

Example Code for Transition to Idle State

To transition to the Idle state, use code shown in Listing 3-3.

Listing 3-3. Transitioning to Idle State

```
CLI R0 ;    /* disable interrupts */
IDLE ;      /* drain pipeline and send core into IDLE state */
STI R0 ;    /* re-enable interrupts after wakeup */
```

Reset State

Reset state initializes the processor logic. During Reset state, application programs and the operating system do not execute. Clocks are stopped while in Reset state.

The processor remains in the Reset state as long as external logic asserts the external $\overline{\text{RESET}}$ signal. Upon de-assertion, the processor completes the reset sequence and switches to Supervisor mode, where it executes code found at the reset event vector.

Software in Supervisor or Emulation mode can invoke the Reset state without involving the external RESET signal. This can be done by issuing the Reset version of the RAISE instruction.

Application programs in User mode cannot invoke the Reset state, except through a system call provided by an operating system kernel. Table 3-5 summarizes the state of the processor upon reset.

Table 3-5. Processor State Upon Reset

Item	Description of Reset State
Cores A and B	
Operating Mode	Supervisor mode in reset event, clocks stopped
Rounding Mode	Unbiased rounding
Cycle Counters	Disabled, zero
DAG Registers (I, L, B, M)	Random values (must be cleared at initialization)
Data and Address Registers	Random values (must be cleared at initialization)
IPEND, IMASK, ILAT	Cleared, interrupts globally disabled with IPEND bit 4
CPLBs	Disabled
L1 Instruction Memory	SRAM (cache disabled)
L1 Data Memory	SRAM (cache disabled)
Cache Validity Bits	Invalid
Core B	Idle
System	
Booting Methods	Determined by the values of BMODE pins at reset
MSEL Clock Frequency	Reset value = 10
PLL Bypass Mode	Disabled
VCO/Core Clock Ratio	Reset value = 1
VCO/System Clock Ratio	Reset value = 5
Peripheral Clocks	Disabled

System Reset and Powerup Configuration

Table 3-6 describes the five types of resets. Note all resets, except System Software, reset the core.

Table 3-6. Resets

Reset	Source	Result
Hardware Reset	The $\overline{\text{RESET}}$ pin causes a hardware reset.	Resets both the core and the peripherals, including the Dynamic Power Management Controller (DPMC). Resets the No Boot on Software Reset bit in SICA_SYSCR. For more information, see "System Reset Configuration Register A (SICA_SYSCR)" on page 3-14.
System Software Reset	Writing b#111 to bits [2:0] in the system MMR SICA_SWRST at address 0xFFC0 0100 causes a System Software reset.	Resets only the peripherals and most of the DPMC. The DPMC resets only the No Boot on Software Reset bit in SYSCR. Does not reset the core. Does not initiate a boot sequence.
Watchdog Timer Reset	Programming the watchdog timer appropriately causes a Watchdog Timer reset.	Resets both the cores and most of the DPMC. The Software Reset register (SICA_SWRST) can be read to determine whether the reset source was the watchdog timer.

Table 3-6. Resets (Cont'd)

Reset	Source	Result
Core Double-Fault Reset	If the core enters a double-fault state, and the Core Double Fault Reset Enable bit (DOUBLE_FAULT) is set in the SWRST register, then a software reset will occur.	Resets both the cores and the peripherals and most of the DPMC. The SICA_SWRST register can be read to determine whether the reset source was Core Double Fault.
Core-Only Software Reset	This reset is caused by executing a RAISE1 instruction or by setting the Software Reset (SYSRST) bit in the core Debug Control register (DBGCTL) via emulation software through the JTAG port. The DBGCTL register is not visible to the memory map.	Resets both the cores. The peripherals do not recognize this reset.

Hardware Reset

The processor chip reset is an asynchronous reset event. The $\overline{\text{RESET}}$ input pin must be de-asserted to perform a hardware reset. For more information, see the applicable *ADSP-BF561 Embedded Processor Data Sheet*.

A hardware-initiated reset results in a system-wide reset that includes both cores and peripherals. After the $\overline{\text{RESET}}$ pin is de-asserted, the processor ensures that all asynchronous peripherals have recognized and completed a reset. After the reset, the processor transitions into the Boot mode sequence configured by the BMODE state.

The BMODE[1:0] pins are dedicated mode control pins. No other functions are shared with these pins, and they may be permanently strapped by tying them directly to either V_{DD} or V_{SS}. The pins and the corresponding bits in SICA_SYSCR configure the Boot mode that is employed after hardware reset or System Software reset. See "Reset" on page 4-41 and Table 4-13 on page 4-44 for further information.

System Reset Configuration Register A (SICA_SYSCR)

The values sensed from the BMODE[1:0] pins are latched into the System Reset Configuration register (SICA_SYSCR) upon the de-assertion of the RESET pin. The values are made available for software access and modification after the hardware reset sequence. Software can modify only the No Boot on Software Reset bit.

The various configuration parameters are distributed to the appropriate destinations from SICA_SYSCR (see Figure 3-2).

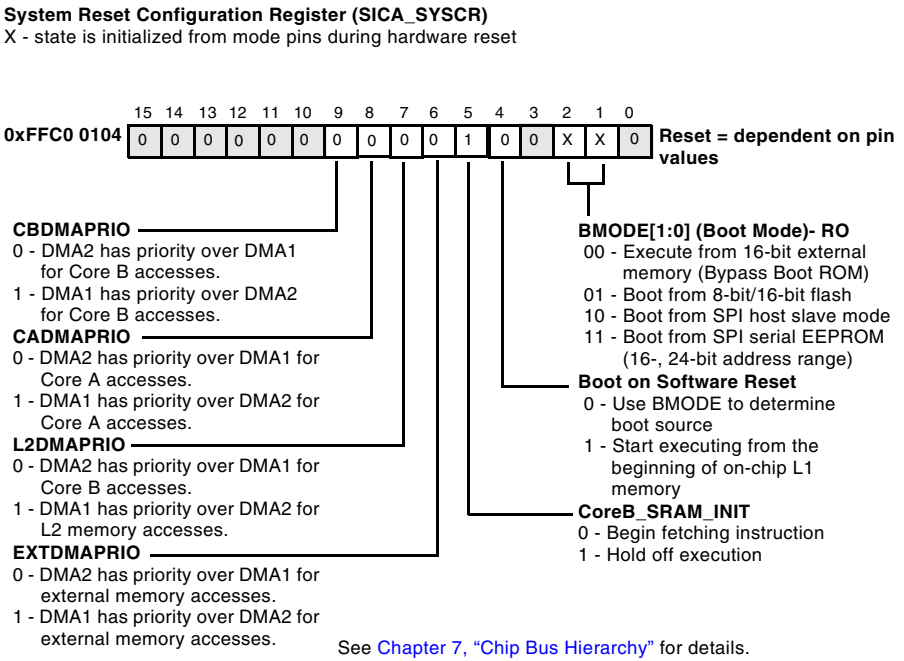

Figure 3-2. System Reset Configuration Register (SICA_SYSCR)

At Reset, core B is “locked”. It remains in this condition until bit 5 of the SICA_SYSCR register is cleared. This can be done explicitly in the program

code, or by the emulator. If core B is not being used, it can remain locked to save power consumption.

System Reset Configuration Register B (SICB_SYSCR)

The SICB_SYSCR register provides the ability for one core to signal the other core via supplemental interrupts. Either core can raise either of the supplemental interrupts on itself or on the other core. The SICB_SYSCR register bits[13:6] are used for this purpose. Figure 3-3 describes this register. Software should write a "1" to bits[9:6] to set the associated supplemental interrupt. Once set, these bits become sticky until the supplemental interrupt is cleared when software writes a "1" to the corresponding bit in bits[13:10].

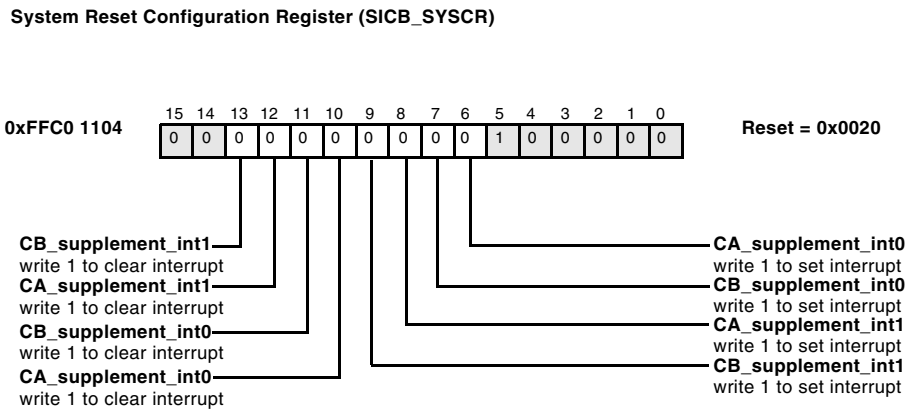

Figure 3-3. System Reset Configuration Register (SICB_SYSCR)

For example, Core A can write a "1" to bit 7 of SICB_SYSCR to interrupt Core B. Core B would then need to write a "1" to bit 11 to clear the interrupt. Similarly, Core B can write a "1" to bit 8 of SICB_SYSCR to interrupt Core A. Core A would then need to write a "1" to bit 12 to clear the interrupt.

This enables either core to interrupt one or both cores at the same time. To enable an interrupt vector as a result of a supplemental interrupt, the supplemental interrupt being used must be enabled in the appropriate SIC_IMASKx register. For more information, see "System Interrupt Mask Registers (SIC_IMASKx)" on page 4-29.

If no interrupt is desired, or if handshaking between the cores is needed for core synchronization reasons, the SICB_SYSCR register can also be used for polling. Since the supplemental interrupt bits[9:6] become sticky once set, reads of this register will always return the status of the supplemental interrupt requests in bits[9:6]. For example, if Core A needs to know when the supplemental interrupt it raises on Core B has been serviced, it can immediately poll the bit it just set and wait for it to clear, provided Core B writes a "1" to the associated clear bit.

Software Resets and Watchdog Timer

A software reset may be initiated in three ways:

- By the watchdog timers, if appropriately configured
- By setting the System Software Reset field in the Software Reset register (see Figure 3-4)
- By the RAISE 1 instruction

The watchdog timer resets both the core and the peripherals. A System Software reset results in a reset of the peripherals without resetting the core and without initiating a booting sequence.

The System Software reset must be performed while executing from Level 1 memory (either as cache or as SRAM).

When L1 instruction memory is configured as cache, make sure the System Software reset sequence has been read into the cache.

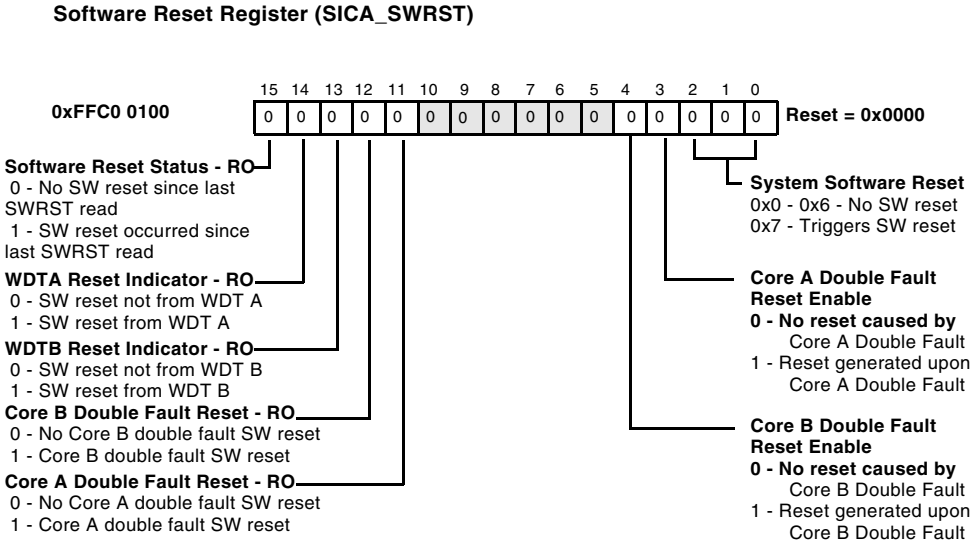

Figure 3-4. Software Reset Register

After either the watchdog or System Software reset is initiated, the processor ensures that all asynchronous peripherals have recognized and completed a reset.

For a reset generated by the watchdog timer, the processors transitions into the Boot mode sequence. The Boot mode is configured by the state of the BMODE and the No Boot on Software Reset control bits.

If the No Boot on Software Reset bit in SYSCR is cleared, the reset sequence is determined by the BMODE[1:0] control bits.

Software Reset Register (SICA_SWRST)

A software reset can be initiated by setting the System Software Reset field in the Software Reset register (SICA_SWRST). Bit 15 indicates whether a software reset has occurred since the last time SICA_SWRST was read. Bit 14 and Bit 13, respectively, indicate whether the Watchdog Timer A or B has generated a software reset. Bit 12 and Bit 11, respectively, indicate

whether the Core A Double Fault or the Core B Double Fault has generated a software reset. Bits [15:11] are read-only and cleared when the register is read. Bits [4:0] are read/write bits.

SICA_SWRST is used to control software reset, while SICB_SWRST is reserved.

When the BMODE pins are not set to b#00, the processor starts executing from the start of on-chip L1 memory. In this configuration, the core begins fetching instructions from the beginning of on-chip L1 memory.

Core-Only Software Reset

A Core-Only Software reset is initiated by executing the RAISE 1 instruction.

A Core-Only Software reset affects only the state of both cores. Note the system resources may be in an undetermined or even unreliable state, depending on the system activity during the reset period.

For dual core environment, any core reset operation will cause dual core reset simultaneously.

Core and System Reset

To perform a system and core reset, use the code sequence shown in Listing 3-4. As described in the code comments below, the system soft reset takes 5 system clock cycles to complete, so a delay loop is needed. This code must reside in L1 memory for the system soft reset to work properly.

Listing 3-4. Core and System Reset

```
/* Issue system soft reset */
P0.L = LO(SICA_SWRST) ;
P0.H = HI(SICA_SWRST) ;
```

```
R0.L = 0x0007 ;
W[P0] = R0 ;
SSYNC ;

/* ******************
Wait for system reset to complete (needs to be 5 SCLKs). Assuming
a worst case CCLK:SCLK ratio (15:1), use 5*15 = 75 as the loop
count.
****************** */
P1 = 75 ;
LSETUP(start, end) LC0 = P1 ;
   start:
   end:
      NOP ;

/* Clear system soft reset */
R0.L = 0x0000 ;
W[P0] = R0 ;
SSYNC ;

/* Core reset - forces reboot */
RAISE 1 ;
```

Booting Methods

The internal boot ROM includes boot kernel code that can either be bypassed or used to load user code from an external memory device. The boot kernel reads the `BMODE[1:0]` pin state at reset to identify the download source. See Table 4-12 on page 4-42 for further information.

Several boot methods are available in which user code can be loaded from an external memory device. For these modes, the boot kernel code sets up the selected peripheral based on the `BMODE[1:0]` pin settings.

At power up or after an hardware reset event, Core A comes out of reset and executes the boot kernel code at address 0xEF00 0000. Core B is held in the idle state.

For boot modes 01 and 11, the boot kernel code reads a boot stream from an external memory device. The boot stream contains user code to be placed at specified memory locations. The boot kernel code will load blocks of user code until it reaches a block that is flagged as the last block. Then it will jump to the start of Core A program memory.

If there are separate executable programs for each core, the Boot ROM will load a second stage boot loader to the start of Core A program memory, then execute the second stage boot loader. The second stage boot loader will then load the separate programs for each core.

Core B will remain in the idle state, until the user code executed by Core A clears the COREB_SRAM_INIT bit in the SICA_SYSCR. When this bit is cleared, Core B is released from idle, and begins executing from the start of its L1 instruction memory space.

Core A can change the PLL settings to set the desired VCO, core and system clock frequencies and then execute an idle instruction. Due to this, both the cores will be in the idle mode and PLL will initiate the desired change in its operating frequencies. After the PLL is locked to the new frequency, it can wake up Core A. Core A can then clear the COREB_SRAM_INIT bit in the SICA_SYSCR register to wake up Core B.

The Core-Only software reset also vectors the Core A to the boot ROM. Only the cores are reset with the Core-Only software reset; this reset does not affect the rest of the system. The boot ROM kernel detects a “No Boot on Software Reset” condition in SICA_SYSCR to avoid initiating a download. If this bit is set on a software reset, Core A skips the normal boot sequence and jumps to the beginning of its L1 memory and begins execution. Core A can then change the PLL settings and/or wake up the core B as described above.

If booting from the Serial Peripheral Interface (SPI), general-purpose flag pin 2 (PF2) is used as the SPI chip select. This line must be connected for proper operation.

The boot kernel assumes these conditions for the Flash Boot mode (BMODE = b#01:

- Asynchronous Memory Bank (AMB) 0 enabled
- 16-bit packing for AMB 0 enabled
- RDY pin not used to determine access completed
- Bank 0 hold time (read/write de-asserted to $\overline{\text{AOE}}$ de-asserted) = 3 cycles
- Bank 0 read/write access times = 15 cycles
- Bank 0 setup time = 4 cycles

For SPI master mode boot (BMODE = 11), the boot kernel assumes that the SPI baud rate is 500 kHz. SPI serial EEPROM and flash memories that are 16- or 24-bit addressable are supported. The SPI uses the PF2 output pin to select a single SPI memory device. For device detection the SPI controller initially submits successive read commands at addresses 0x00, 0x0000, and 0x000000 until a valid 8-, 16- or 24-bit addressable memory device is detected. It then begins reading and processing the boot stream using the proper addressing format.

SPI Slave Mode Boot Through Master Host (BMODE = 10)

For SPI slave mode booting, the ADSP-BF561 is configured as an SPI slave device and a host is used to boot the processor. The host drives the SPI clock and is therefore responsible for the timing. The baud rate should be equal to or less than one fourth (1/4) of the ADSP-BF561 system clock (SCLK).

This boot mode is only supported starting with silicon revision 0.5.

The host must ensure that the Blackfin processor is out of reset before sending any bytes. Any bytes sent to the Blackfin processor prior to reset will be lost and the booting sequence will fail.

The host does not need any knowledge of the loader file stream to boot the Blackfin processor. The host must be configured to send one byte at a time from the loader file (ASCII format). A programmable flag `PFx` must be used as the host wait (`HWAIT`) signal from the Blackfin processor to the master device. This is the signal used by the Blackfin processor to *hold off* the host during certain times within the boot process (specifically during initialization code execution and zero-fill blocks). When `PFx` is asserted (high), the master host device must discontinue sending bytes to the Blackfin processor. When `PFx` is de-asserted (low), the master host resumes sending bytes from where it left off. Since the `PFx` pin is not driven by the slave until the first block has been processed, consider using a resistor to pull down the `HWAIT` signal.

This `PFx` number is user-defined and embedded within the loader file. Valid values for `HWAIT` are 1 through 15. This can be selected either through the Loader Property page in VisualDSP++ IDDE, or through the command line by using the `-pflag` switch (for example, `-pflag 10`) to select `PF10` as the `HWAIT` signal.

4 PROGRAM SEQUENCER

The processor has two identical cores and each core has its own program sequencer. The Program Sequencer controls program flow, constantly providing the address of the next instruction to be executed by other parts of the processor. Program flow in the chip is mostly linear, with the processor executing program instructions sequentially.

The linear flow varies occasionally when the program uses nonsequential program structures, such as those illustrated in Figure 4-1. Nonsequential structures direct the processor to execute an instruction that is not at the next sequential address. These structures include:

- **Loops.** One sequence of instructions executes several times with zero overhead.
- **Subroutines.** The processor temporarily interrupts sequential flow to execute instructions from another part of memory.
- **Jumps.** Program flow transfers permanently to another part of memory.
- **Interrupts and Exceptions.** A runtime event or instruction triggers the execution of a subroutine.
- **Idle**. An instruction causes the processor to stop operating and hold its current state until an interrupt occurs. Then, the processor services the interrupt and continues normal execution.

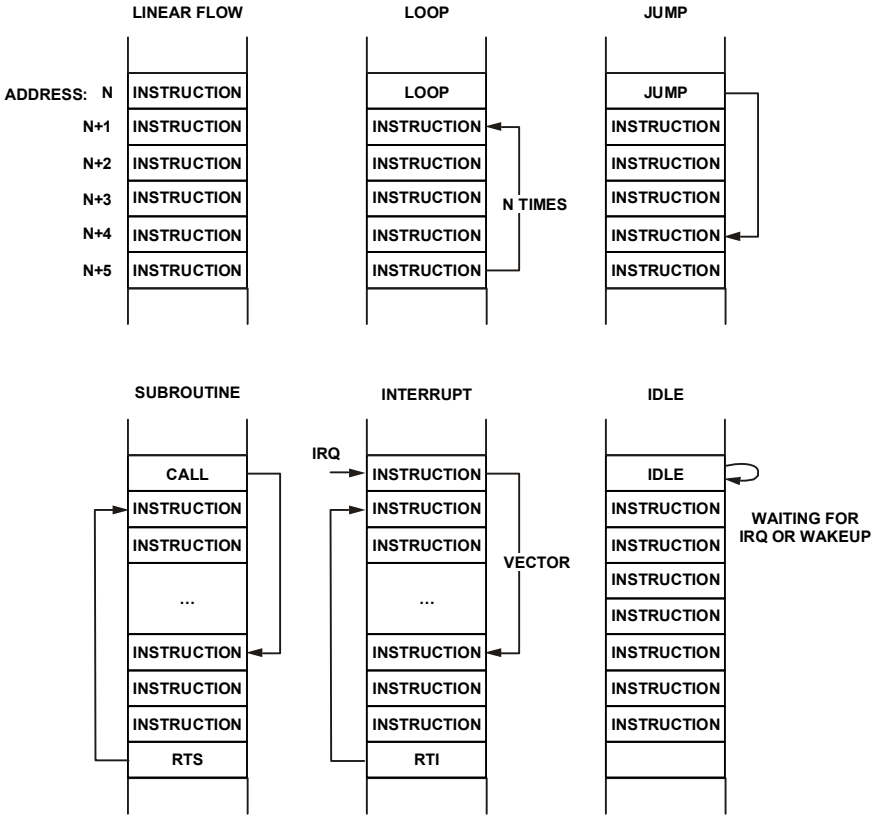

Figure 4-1. Program Flow Variations

The Sequencer manages execution of these program structures by selecting the address of the next instruction to execute.

The fetched address enters the instruction pipeline, ending with the program counter (PC). The pipeline contains the 32-bit addresses of the instructions currently being fetched, decoded, and executed. The PC couples with the RETn registers, which store return addresses. All addresses generated by the Sequencer are 32-bit memory instruction addresses.

To manage events, the Sequencer's event controller handles interrupt and event processing, determines whether an interrupt is masked, and generates the appropriate event vector address.

In addition to providing data addresses, the Data Address Generators (DAGs) can provide instruction addresses for the Sequencer's indirect branches.

The Sequencer evaluates conditional instructions and loop termination conditions. The loop registers support nested loops. The memory-mapped registers (MMRs) store information used to implement interrupt service routines.

Sequencer-Related Registers

Table 4-1 lists the registers within the processor that are related to each core's Sequencer. Both cores have the same Sequencer-related register set. Except for the `PC` and `SEQSTAT` registers, all Sequencer-related registers are directly readable and writable. Manually pushing or popping registers to or from the stack is done using the explicit instructions `[--SP] = Rn` (for push) or `Rn = [SP++]` (for pop).

Table 4-1. Sequencer-Related Registers for Each Core

Register Name	Description
`SEQSTAT`	Sequencer Status register

`RETX`
`RETN`
`RETI`
`RETE`
`RETS`
	Return Address registers: (See "Events and Sequencing" on page 4-18.)
	Exception Return
	NMI Return
	Interrupt Return
	Emulation Return
	Subroutine Return

Table 4-1. Sequencer-Related Registers for Each Core (Cont'd)

Register Name	Description
LC0, LC1 LT0, LT1 LB0, LB1	Zero-Overhead Loop registers: Loop Counters Loop Tops Loop Bottoms
FP, SP	Frame Pointer and Stack Pointer. See Chapter 5, "Data Address Generators".
SYSCFG	System Configuration Register
CYCLES, CYCLES2	Cycle Counters. See Chapter 18, "Blackfin Processor Debug".
PC	Program Counter

Sequencer Status Register (SEQSTAT)

The Sequencer Status register (SEQSTAT) shown in Figure 4-2 contains information about the current state of the Sequencer as well as diagnostic information from the last event. SEQSTAT is a read-only register and is accessible only in Supervisor mode.

Zero-Overhead Loop Registers (LC, LT, LB)

Two sets of zero-overhead loop registers implement loops, using hardware counters instead of software instructions to evaluate loop conditions. After evaluation, processing branches to a new target address. Both sets of registers include the Loop Counter (LC), Loop Top (LT), and Loop Bottom (LB) registers.

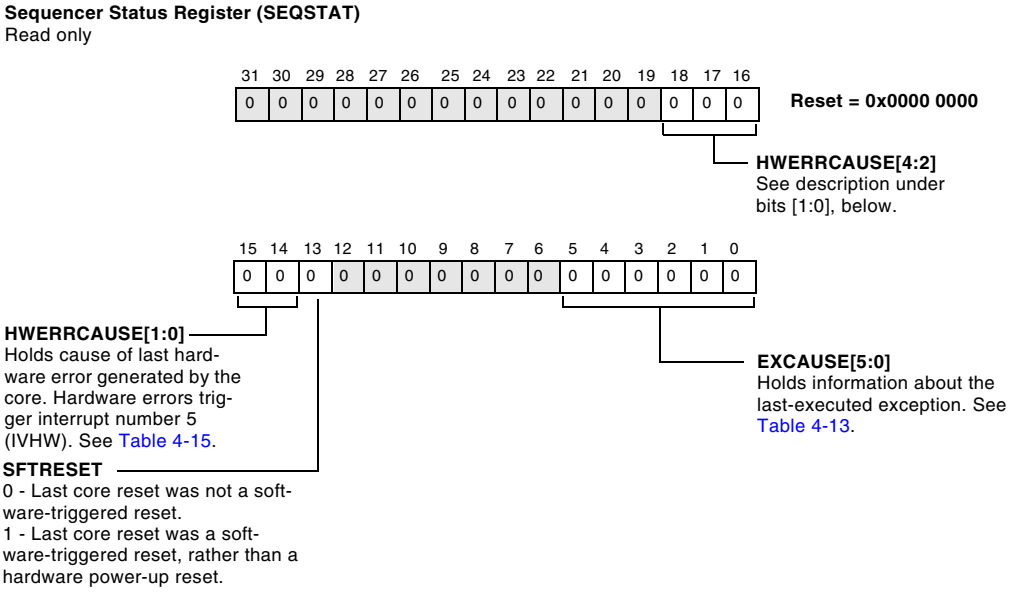

Figure 4-2. Sequencer Status Register

Table 4-2 describes the 32-bit loop register sets.

Table 4-2. Loop Registers

Registers	Description	Function
LC0, LC1	Loop Counters	Maintain a count of the remaining iterations of the loop
LT0, LT1	Loop Tops	Hold the address of the first instruction within a loop
LB0, LB1	Loop Bottoms	Hold the address of the last instruction of the loop

System Configuration Register (SYSCFG)

The System Configuration Register (SYSCFG) shown in Figure 4-3 controls the configuration of the processor. This register is accessible only from the Supervisor mode.

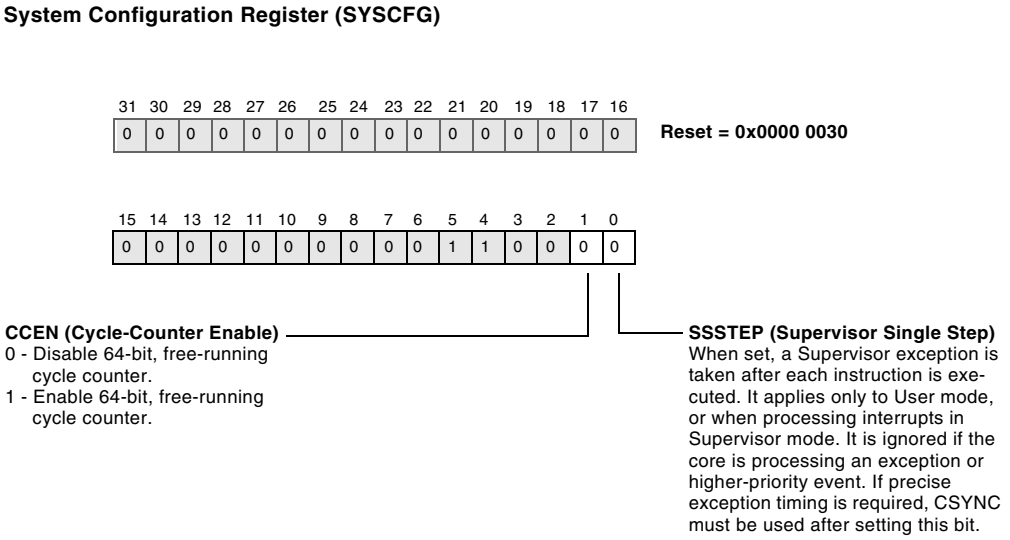

Figure 4-3. System Configuration Register

Instruction Pipeline

The program sequencer determines the next instruction address by examining both the current instruction being executed and the current state of the processor. If no conditions require otherwise, the processor executes instructions from memory in sequential order by incrementing the look-ahead address.

The processor has a ten-stage instruction pipeline, shown in Table 4-3.

Table 4-3. Stages of Instruction Pipeline

Pipeline Stage	Description
Instruction Fetch 1 (IF1)	Issue instruction address to IAB bus, start compare tag of instruction cache
Instruction Fetch 2 (IF2)	Wait for instruction data
Instruction Fetch 3 (IF3)	Read from IDB bus and align instruction
Instruction Decode (DEC)	Decode instructions
Address Calculation (AC)	Calculation of data addresses and branch target address
Data Fetch 1 (DF1)	Issue data address to DA0 and DA1 bus, start compare tag of data cache
Data Fetch 2 (DF2)	Read register files
Execute 1 (EX1)	Read data from LD0 and LD1 bus, start multiply and video instructions
Execute 2 (EX2)	Execute/Complete instructions (shift, add, logic, etc.)
Write Back (WB)	Writes back to register files, SD bus, and pointer updates (also referred to as the “commit” stage)

Figure 4-4 shows a diagram of the pipeline.

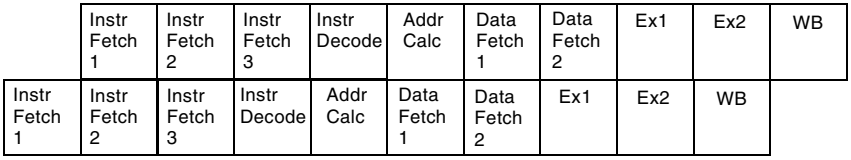

Figure 4-4. Processor Pipeline

The instruction fetch and branch logic generates 32-bit fetch addresses for the Instruction Memory Unit. The Instruction Alignment Unit returns instructions and their width information at the end of the IF3 stage.

For each instruction type (16, 32, or 64 bits), the Instruction Alignment Unit ensures that the alignment buffers have enough valid instructions to be able to provide an instruction every cycle. Since the instructions can be 16, 32, or 64 bits wide, the Instruction Alignment Unit may not need to fetch an instruction from the cache every cycle. For example, for a series of 16-bit instructions, the Instruction Alignment Unit gets an instruction from the Instruction Memory Unit once in four cycles. The alignment logic requests the next instruction address based on the status of the alignment buffers. The sequencer responds by generating the next fetch address in the next cycle, provided there is no change of flow.

The sequencer holds the fetch address until it receives a request from the alignment logic or until a change of flow occurs. The sequencer always increments the previous fetch address by 8 (the next 8 bytes). If a change of flow occurs, such as a branch or an interrupt, data in the Instruction Alignment Unit is invalidated. The sequencer decodes and distributes instruction data to the appropriate locations such as the register file and data memory.

The Execution Unit contains two 16-bit multipliers, two 40-bit ALUs, two 40-bit accumulators, one 40-bit shifter, a video unit (which adds 8-bit ALU support), and an 8-entry 32-bit Data Register File.

Register file reads occur in the DF2 pipeline stage (for operands).

Register file writes occur in the WB stage (for stores). The multipliers and the video units are active in the EX1 stage, and the ALUs and shifter are active in the EX2 stage. The accumulators are written at the end of the EX2 stage.

The program sequencer also controls stalling and invalidating the instructions in the pipeline. Multi-cycle instruction stalls occur between the IF3 and DEC stages. DAG and sequencer stalls occur between the DEC and AC stages. Computation and register file stalls occur between the DF2 and EX1 stages. Data memory stalls occur between the EX1 and EX2 stages.

The sequencer ensures that the pipeline is fully interlocked and that all the data hazards are hidden from the programmer.

Multi-cycle instructions behave as multiple single-cycle instructions being issued from the decoder over several clock cycles. For example, the Push Multiple or Pop Multiple instruction can push or pop from 1 to 14 DREGS and/or PREGS, and the instruction remains in the decode stage for a number of clock cycles equal to the number of registers being accessed.

Multi-issue instructions are 64 bits in length and consist of one 32-bit instruction and two 16-bit instructions. All three instructions execute in the same amount of time as the slowest of the three.

Any nonsequential program flow can potentially decrease the processor's instruction throughput. Nonsequential program operations include:

- Jumps
- Subroutine calls and returns
- Interrupts and returns
- Loops

Branches and Sequencing

One type of nonsequential program flow that the Sequencer supports is branching. A branch occurs when a JUMP or CALL instruction begins execution at a new location other than the next sequential address. For descriptions of how to use the JUMP and CALL instructions, see the *Blackfin Processor Programming Reference*. Briefly, these instructions operate as follows:

- A JUMP or a CALL instruction transfers program flow to another memory location. The difference between a JUMP and a CALL is that a CALL automatically loads the return address into the RETS register.

The return address is the next sequential address after the CALL instruction. This push makes the address available for the CALL instruction's matching return instruction, allowing easy return from the subroutine.

- A return instruction causes the Sequencer to fetch the instruction at the return address, which is stored in the RETS register (for subroutine returns). The types of return instructions are return from subroutine (RTS), return from interrupt (RTI), return from exception (RTX), return from emulation (RTE), and return from nonmaskable interrupt (RTN). Each return type has its own register for holding the return address.
- JUMP instructions can be conditional, depending on the status of the CC bit of the ASTAT register. They are immediate and may not be delayed. The Program Sequencer can evaluate the CC status bit to decide whether to execute a branch. If no condition is specified, the branch is always taken.
- Conditional JUMP instructions use static branch prediction to reduce the branch latency caused by the length of the pipeline.

Branches can be direct or indirect. A direct branch address is determined solely by the instruction word (for example JUMP 0x30), while an indirect branch gets its address from the contents of a DAG register (for example, JUMP (P3)).

All types of JUMPS and CALLS can be PC-relative. The indirect JUMP and CALL can be absolute or PC-relative.

Direct Short and Long Jumps

The Sequencer supports both short and long jumps. The target of the branch is a PC-relative address from the location of the instruction, plus an offset. The PC-relative offset for the short jump is a 13-bit immediate

value that must be a multiple of two (bit zero must be a zero). The 13-bit value gives an effective dynamic range of -4096 to +4094 bytes.

The PC-relative offset for the long jump is a 25-bit immediate value that must also be a multiple of two (bit zero must be a zero). The 25-bit value gives an effective dynamic range of –16,777,216 to +16,777,214 bytes.

If, at the time of writing the program, the destination is known to be less than a 13-bit offset from the current PC value, then the `JUMP.S 0xnnnn` instruction may be used. If the destination requires more than a 13-bit offset, then the `JUMP.L 0xnnnn nnn` instruction must be used. If the destination offset is unknown and development tools must evaluate the offset, then use the instruction `JUMP 0xnnnn nnn`. Upon disassembly, the instruction is replaced by the appropriate `JUMP.S` or `JUMP.L` instruction.

Direct Call

The `CALL` instruction is a branch instruction that copies the address of the instruction which would have executed next (had the `CALL` instruction not executed) into the `RETS` register. The direct `CALL` instruction has a 25-bit PC-relative offset that must be a multiple of two (bit zero must be a zero). The 25-bit value gives an effective dynamic range of –16,777,216 to +16,777,214 bytes.

Indirect Branch and Call

The indirect `JUMP` and `CALL` instructions get their destination address from a data address generator (DAG) P-register. For the `CALL` instruction, the `RETS` register is loaded with the address of the instruction which would have executed next in the absence of the `CALL` instruction.

For example:

```
JUMP (P3);
CALL (P0);
```

PC-Relative Indirect Branch and Call

The PC-relative indirect JUMP and CALL instructions use the contents of a P-register as an offset to the branch target. For the CALL instruction, the RETS register is loaded with the address of the instruction which would have executed next (had the CALL instruction not executed).

For example:

```
JUMP (PC + P3);
CALL (PC + P0);
```

Condition Code Flag

The processor supports a condition code (CC) flag bit, which is used to resolve the direction of a branch. This flag may be accessed eight ways:

- A conditional branch is resolved by the value in CC.
- A Data register value may be copied into CC, and the value in CC may be copied to a Data register.
- The BITTST instruction accesses the CC flag.
- A status flag may be copied into CC, and the value in CC may be copied to a status flag.
- CC may be set to the result of a Pointer register comparison.
- CC bit may be set to the result of a Data register comparison.
- Some shifter instructions (rotate or BXOR) use CC as a portion of the shift operand/result.
- Test and set instructions can set and clear the CC bit.

These eight ways of accessing the CC bit are used to control program flow. The branch is explicitly separated from the instruction that sets the

arithmetic flags. A single bit resides in the instruction encoding that specifies the interpretation for the value of CC. The interpretation is to "branch on true" or "branch on false."

The comparison operations have the form CC = expr where *expr* involves a pair of registers of the same type (for example, Data registers or Pointer registers, or a single register and a small immediate constant). The small immediate constant is a 3-bit (-4 through 3) signed number for signed comparisons and a 3-bit (0 through 7) unsigned number for unsigned comparisons.

The sense of CC is determined by equal (==), less than (<), and less than or equal to (<=). There are also bit test operations that test whether a bit in a 32-bit register is set.

Conditional Branches

The Sequencer supports conditional branches. These are JUMP instructions whose execution branches or continues linearly, depending on the value of the CC bit. The target of the branch is a PC-relative address from the location of the instruction, plus an offset. The PC-relative offset is an 11-bit immediate value that must be a multiple of two (bit zero must be a zero). This gives an effective dynamic range of –1024 to +1022 bytes.

For example, the following instruction tests the CC flag and, if it is positive, jumps to a location identified by the label dest_address:

```
IF CC JUMP dest_address;
```

Conditional Register Move

Register moves can be performed depending on whether the value of the CC flag is true or false (1 or 0). In some cases, using this instruction instead of a branch eliminates the cycles lost because of the branch. These

conditional moves can be done between any R- or P-registers (including SP and FP).

Example code:

```
IF CC R0 = P0;
```

Branch Prediction

The Sequencer supports static branch prediction to accelerate execution of conditional branches. These branches are executed based on the state of the CC bit.

In the EX4 stage, the Sequencer compares the actual CC bit value to the predicted value. If the value was mis-predicted, the branch is corrected, and the correct address is available for the WB stage of the pipeline.

The branch latency for conditional branches is as follows:

- If prediction was “not to take branch,” and branch was actually not taken: 0 CCLK cycles.
- If prediction was “not to take branch,” and branch was actually taken: 8 CCLK cycles.
- If prediction was “to take branch,” and branch was actually taken: 4 CCLK cycles.
- If prediction was “to take branch,” and branch was actually not taken: 8 CCLK cycles.

For all unconditional branches, the branch target address computed in the AC stage of the pipeline is sent to the Instruction Fetch address bus at the beginning of the EX1 stage. All unconditional branches have a latency of 4 CCLK cycles.

Consider the example in Table 4-4.

Table 4-4. Branch Prediction

Instruction	Description
`If CC JUMP dest (bp)`	This instruction tests the `CC` flag, and if it is set, jumps to a location identified by the label, `dest`. If the `CC` flag is set, the branch is correctly predicted and the branch latency is reduced. Otherwise, the branch is incorrectly predicted and the branch latency increases.

Loops and Sequencing

The sequencer supports a mechanism of zero-overhead looping. The sequencer contains two loop units, each containing three registers. Each loop unit has a Loop Top register (`LT0`, `LT1`), a Loop Bottom register (`LB0`, `LB1`), and a Loop Count register (`LC0`, `LC1`).

When an instruction at address `X` is executed, and `X` matches the contents of `LB0`, then the next instruction executed will be from the address in `LT0`. In other words, when `PC == LB0`, then an implicit jump to `LT0` is executed.

A loopback only occurs when the count is greater than or equal to 2. If the count is nonzero, then the count is decremented by 1. For example, consider the case of a loop with two iterations. At the beginning, the count is 2. Upon reaching the first loop end, the count is decremented to 1 and the program flow jumps back to the top of the loop (to execute a second time). Upon reaching the end of the loop again, the count is decremented to 0, but no loopback occurs (because the body of the loop has already been executed twice).

Since there are two loop units, loop unit 1 is assigned higher priority so it can be used as the inner loop in a nested loop structure. In other words, a loopback caused by loop unit 1 on a particular instruction (`PC == LB1`, `LC1 >= 2`) will prevent loop unit 0 from looping back on that same

instruction, even if the address matches. Loop unit 0 is allowed to loop back only after the loop count 1 is exhausted.

The LSETUP instruction can be used to load all three registers of a loop unit at once. Each loop register can also be loaded individually with a register transfer, but this incurs a significant overhead if the loop count is nonzero (the loop is active) at the time of the transfer.

The code example in Listing 4-1 shows a loop that contains two instructions and iterates 32 times.

Listing 4-1. Loop Example

```
P5 = 0x20;
LSETUP (lp_start, lp_end) LC0 = P5;
lp_start:
R5 = R0 + R1(ns) || R2 = [P2++] || R3 = [I1++];

lp_end: R5 = R5 + R2;
```

Two sets of loop registers are used to manage two nested loops:

- LC[1:0] – the Loop Count registers
- LT[1:0] – the Loop Top address registers
- LB[1:0] – the Loop Bottom address registers

When executing an LSETUP instruction, the Program Sequencer loads the address of the loop's last instruction into LBx and the address of the loop's first instruction into LTx. The top and bottom addresses of the loop are computed as PC-relative addresses from the LSETUP instruction plus an offset. In each case, the offset value is added to the location of the LSETUP instruction.

LC0 and LC1 are unsigned 32-bit registers, each supporting 2^{32} -1 iterations through the loop.

When LCx = 0, the loop is disabled, and a single pass of the code executes.

Table 4-5. Loop Registers

First/Last Address of the Loop	PC-Relative Offset Used to Compute the Loop Start Address	Effective Range of the Loop Start Instruction
Top / First	5-bit signed immediate; must be a multiple of 2.	0 to 30 bytes away from LSETUP instruction.
Bottom / Last	11-bit signed immediate; must be a multiple of 2.	0 to 2046 bytes away from LSETUP instruction (the defined loop can be 2046 bytes long).

The processor supports a four-location instruction loop buffer that reduces instruction fetches while in loops. If the loop code contains four or fewer instructions, then no fetches to instruction memory are necessary for any number of loop iterations, because the instructions are stored locally. The loop buffer effectively eliminates the instruction fetch time in loops with more than four instructions by allowing fetches to take place while instructions in the loop buffer are being executed.

A four-cycle latency occurs on the first loopback when the LSETUP specifies a non-zero start offset (lp_start). Therefore, zero start offsets are preferred.

The processor has no restrictions regarding which instructions can occur in a loop end position. Branches and calls are both allowed in that position.

Events and Sequencing

The Event Controller of each core manages five types of activities:

- Emulation
- Reset
- Non-maskable interrupts (NMI)
- Exceptions
- Interrupts

Note that the word *event* describes all five types. The Event Controller manages fifteen events in all: Emulation, Reset, NMI, Exception, and eleven Interrupts.

An interrupt is an event that changes normal processor instruction flow and is asynchronous to program flow. In contrast, an exception is a software-initiated event whose effects are synchronous to program flow.

The event system is nested and prioritized. Consequently, several service routines may be active at any time, and a low-priority event may be preempted by one of higher priority.

Each processor core employs a two-level event control mechanism. The System Interrupt Controller (SIC) works with the Core Event Controller (CEC) to prioritize and control all system interrupts. The SIC provides mapping between the many peripheral interrupt sources and the prioritized general-purpose interrupt inputs of the core. This mapping is programmable, and individual interrupt sources can be masked in the SIC.

The CEC supports nine general-purpose interrupts (`IVG7 - IVG15`) in addition to the dedicated interrupt and exception events that are described in Table 4-6. It is recommended that the two lowest priority interrupts (`IVG14` and `IVG15`) be reserved for software interrupt handlers, leaving

seven prioritized interrupt inputs (IVG7 - IVG13) to support the system. Refer to Table 4-6 for more information.

Table 4-6. System and Core Event Mapping

	Event Source	Core Event Name
Core Events	Emulation (highest priority)	EMU
	Reset	RST
	NMI	NMI
	Exception	EVX
	Reserved	–
	Hardware Error	IVHW
	Core Timer	IVTMR
System Interrupts (See Table 4-7)	General Interrupt 7	IVG7
	General Interrupt 8	IVG8
	General Interrupt 9	IVG9
	General Interrupt 10	IVG10
	General Interrupt 11	IVG11
	General Interrupt 12	IVG12
	General Interrupt 13	IVG13
	Software Interrupt 1	IVG14
	Software Interrupt 2 (lowest priority)	IVG15

Note that the System Interrupt to Core Event mappings shown in Table 4-6 are the default values at reset and can be changed by software.

System Interrupt Processing

Processor Core A and Core B have separate system interrupt registers (for example, SICA_IWR0 and SICB_IWR0). For clarity, this chapter instead refers to the generic system interrupt register (for example, SIC_IWR0). Please refer to Appendix B, "System MMR Assignments" for the core-specific MMR names and addresses.

By setting the appropriate system interrupt registers, a peripheral interrupt source can generate an interrupt or a wakeup to either core, or to both cores.

Referring to Figure 4-5, note that when an interrupt (Interrupt A) is generated by an interrupt-enabled peripheral:

1. SIC_ISRx logs the request and keeps track of system interrupts that are asserted but not yet serviced (that is, an interrupt service routine hasn't yet cleared the interrupt).
2. SIC_IWRx checks to see if it should wake up the core from an idled state based on this interrupt request.
3. SIC_IMASKx masks off or enables interrupts from peripherals at the system level. If Interrupt A is not masked, the request proceeds to Step 4.
4. The SIC_IARx registers, which map the peripheral interrupts to a smaller set of general-purpose core interrupts (IVG7-IVG15), determine the core priority of Interrupt A.
5. ILAT adds Interrupt A to its log of interrupts latched by the core but not yet actively being serviced.
6. IMASK masks off or enables events of different core priorities. If the IVGx event corresponding to Interrupt A is not masked, the process proceeds to Step 7.

7. The Event Vector Table (EVT) is accessed to look up the appropriate vector for Interrupt A's interrupt service routine (ISR).
8. When the event vector for Interrupt A has entered the core pipeline, the appropriate IPEND bit is set, which clears the respective ILAT bit. Thus, IPEND tracks all pending interrupts, as well as those being presently serviced.
9. When the interrupt service routine (ISR) for Interrupt A has been executed, the RTI instruction clears the appropriate IPEND bit. However, the relevant SIC_ISRx bit is not cleared unless the interrupt service routine clears the mechanism that generated Interrupt A, or if the process of servicing the interrupt clears this bit.

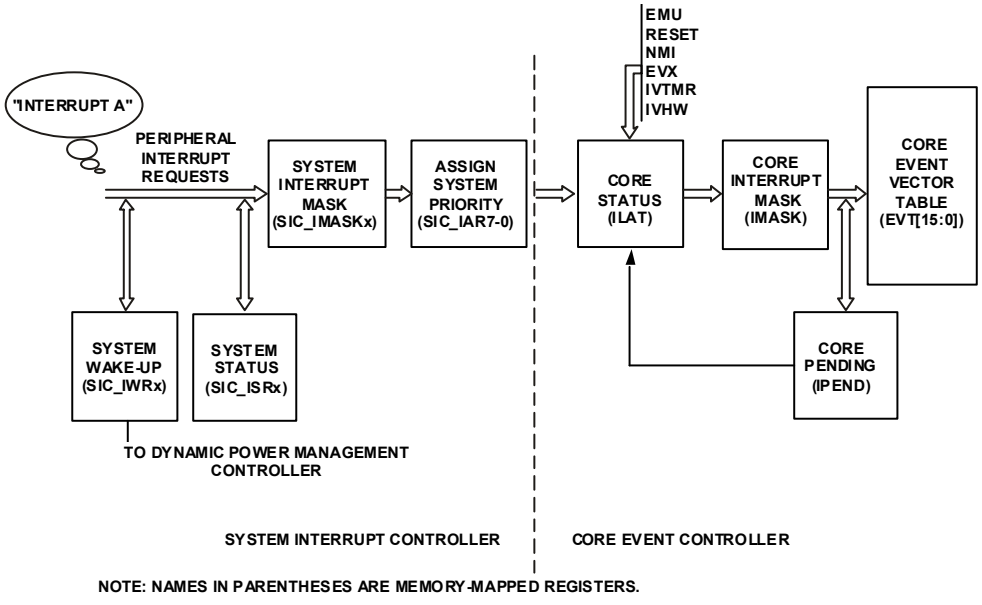

Figure 4-5. Interrupt Processing Block Diagram

It should be noted that emulation, reset, NMI, and exception events, as well as Hardware Error (IVHW) and Core Timer (IVTMR) interrupt requests, enter the interrupt processing chain at the ILAT level and are not affected by the system-level interrupt registers (SIC_IWRx, SIC_ISRx, SIC_IMASKx, SIC_IARx).

If multiple interrupt sources share a single core interrupt, then the ISR must identify the peripheral that generated the interrupt. The ISR may then need to interrogate the peripheral to determine the appropriate action to take.

System Peripheral Interrupts

The processor system has numerous peripherals, which therefore require many supporting interrupts. Table 4-7 lists:

- Peripheral interrupt source.
- Peripheral interrupt ID used in the System Interrupt Assignment registers (SIC_IARx). See "System Interrupt Assignment Registers (SIC_IARx)" on page 4-30.
- General-purpose interrupt of the core to which the interrupt maps at reset.

Table 4-7. Peripheral Interrupt Source Reset State

Peripheral Interrupt Source	Peripheral Interrupt Channel	General Purpose Interrupt (Assignment at Reset)	SIC_IWR0, SIC_ISR0, SIC_IMASK0	SIC_IWR1, SIC_ISR1, SIC_IMASK1
PLL Wakeup	0	IVG7	0	-
DMA1 Error (generic)	1	IVG7	1	-
DMA2 Error (generic)	2	IVG7	2	-
IMDMA Error	3	IVG7	3	-

Table 4-7. Peripheral Interrupt Source Reset State (Cont'd)

Peripheral Interrupt Source	Peripheral Interrupt Channel	General Purpose Interrupt (Assignment at Reset)	SIC_IWR0, SIC_ISR0, SIC_IMASK0	SIC_IWR1, SIC_ISR1, SIC_IMASK1
PPI0 Error	4	IVG7	4	-
PPI1 Error	5	IVG7	5	-
SPORT0 Error	6	IVG7	6	-
SPORT1 Error	7	IVG7	7	-
SPI Error	8	IVG7	8	-
UART Error	9	IVG7	9	-
Reserved	10	IVG7	10	-
DMA1 Channel 0 interrupt (PPI0)	11	IVG8	11	-
DMA1 Channel 1 interrupt (PPI1)	12	IVG8	12	-
DMA1 Channel 2 interrupt	13	IVG8	13	-
DMA1 Channel 3 interrupt	14	IVG8	14	-
DMA1 Channel 4 interrupt	15	IVG8	15	-
DMA1 Channel 5 interrupt	16	IVG8	16	-
DMA1 Channel 6 interrupt	17	IVG8	17	-
DMA1 Channel 7 interrupt	18	IVG8	18	-
DMA1 Channel 8 interrupt	19	IVG8	19	-
DMA1 Channel 9 interrupt	20	IVG8	20	-
DMA1 Channel 10 interrupt	21	IVG8	21	-
DMA1 Channel 11 interrupt	22	IVG8	22	-
DMA2 Channel 0 interrupt (SPORT0 RX)	23	IVG9	23	-
DMA2 Channel 1 interrupt (SPORT0 TX)	24	IVG9	24	-
DMA2 Channel 2 interrupt (SPORT1 RX)	25	IVG9	25	-
DMA2 Channel 3 interrupt (SPORT1 TX)	26	IVG9	26	-

Table 4-7. Peripheral Interrupt Source Reset State (Cont'd)

Peripheral Interrupt Source	Peripheral Interrupt Channel	General Purpose Interrupt (Assignment at Reset)	SIC_IWR0, SIC_ISR0, SIC_IMASK0	SIC_IWR1, SIC_ISR1, SIC_IMASK1
DMA2 Channel 4 interrupt (SPI)	27	IVG9	27	-
DMA2 Channel 5 interrupt (UART RX)	28	IVG9	28	-
DMA2 Channel 6 interrupt (UART TX)	29	IVG9	29	-
DMA2 Channel 7 interrupt	30	IVG9	30	-
DMA2 Channel 8 interrupt	31	IVG9	31	-
DMA2 Channel 9 interrupt	32	IVG9	-	0
DMA2 Channel 10 interrupt	33	IVG9	-	1
DMA2 Channel 11 interrupt	34	IVG9	-	2
Timer0 interrupt	35	IVG10	-	3
Timer1 interrupt	36	IVG10	-	4
Timer2 interrupt	37	IVG10	-	5
Timer3 interrupt	38	IVG10	-	6
Timer4 interrupt	39	IVG10	-	7
Timer5 interrupt	40	IVG10	-	8
Timer6 interrupt	41	IVG10	-	9
Timer7 interrupt	42	IVG10	-	10
Timer8 interrupt	43	IVG10	-	11
Timer9 interrupt	44	IVG10	-	12
Timer10 interrupt	45	IVG10	-	13
Timer11 interrupt	46	IVG10	-	14
Programmable Flags 15-0 interrupt A	47	IVG11	-	15
Programmable Flags 15-0 interrupt B	48	IVG11	-	16
Programmable Flags 31-16 interrupt A	49	IVG11	-	17

Table 4-7. Peripheral Interrupt Source Reset State (Cont'd)

Peripheral Interrupt Source	Peripheral Interrupt Channel	General Purpose Interrupt (Assignment at Reset)	SIC_IWR0, SIC_ISR0, SIC_IMASK0	SIC_IWR1, SIC_ISR1, SIC_IMASK1
Programmable Flags 31-16 interrupt B	50	IVG11	-	18
Programmable Flags 47-32 interrupt A	51	IVG11	-	19
Programmable Flags 47-32 interrupt B	52	IVG11	-	20
DMA1 Channel 12/13 interrupt (Memory DMA Stream 0)	53	IVG8	-	21
DMA1 Channel 14/15 interrupt (Memory DMA Stream 1)	54	IVG8	-	22
DMA2 Channel 12/13 interrupt (Memory DMA Stream 0)	55	IVG9	-	23
DMA2 Channel 14/15 interrupt (Memory DMA Stream 1)	56	IVG9	-	24
IMDMA Stream 0 interrupt	57	IVG12	-	25
IMDMA Stream 1 interrupt	58	IVG12	-	26
Watchdog Timer interrupt	59	IVG13	-	27
Reserved	60	IVG7	-	28
Reserved	61	IVG7	-	29
Supplemental interrupt 0	62	IVG7	-	30
Supplemental interrupt 1	63	IVG7	-	31

The peripheral interrupt structure of the ADSP-BF561 is flexible. By default upon reset, multiple peripheral interrupts share a single, general-purpose interrupt in the core, as shown in Table 4-7.

An interrupt service routine that supports multiple interrupt sources must interrogate the appropriate system MMRs to determine which peripheral generated the interrupt.

If the default assignments shown in Table 4-7 are acceptable, then interrupt initialization involves only initialization of the core EVT vector address entries and IMASK register, and unmasking the specific peripheral interrupts in SIC_IMASK that the system requires.

System Interrupt Wakeup-Enable Registers (SIC_IWRx)

The SIC_IWRx registers provide the mapping between the peripheral interrupt source and the Dynamic Power Management Controller (DPMC). Any of the peripherals can be configured to wake up the core from its idled state to process the interrupt, simply by enabling the appropriate bit in the System Interrupt Wakeup-Enable registers SIC_IWR0 or SIC_IWR1 (refer to Figure 4-6). If a peripheral interrupt source is enabled in SIC_IWRx and the core is idled, the interrupt causes the DPMC to initiate the core wake-up sequence in order to process the interrupt. Note that this mode of operation may add latency to interrupt processing, depending on the power control state. For further discussion of power modes and the idled state of the core, see Chapter 8, "Dynamic Power Management".

By default, all interrupts generate a wake-up request to the core. However, for some applications it may be desirable to disable this function for some peripherals, such as for a SPORTx Transmit Interrupt.

The SIC_IWR registers have no effect unless the core is idled. The bits in this register correspond to those of the System Interrupt Mask (SIC_IMASKx) and Interrupt Status (SIC_ISRx) registers.

After reset, all valid bits of these registers are set to 1, enabling the wake-up function for all interrupts that are not masked. Before enabling interrupts, configure these registers in the reset initialization sequence. SIC_IWRx can be read from or written to at any time. To prevent spurious or lost interrupt activity, these registers should be written only when all peripheral interrupts are disabled.

Note that the wake-up function is independent of the interrupt mask function. If an interrupt source is enabled in SIC_IWRx but masked off in SIC_IMASKx, the core wakes up if it is idled, but it does not generate an interrupt.

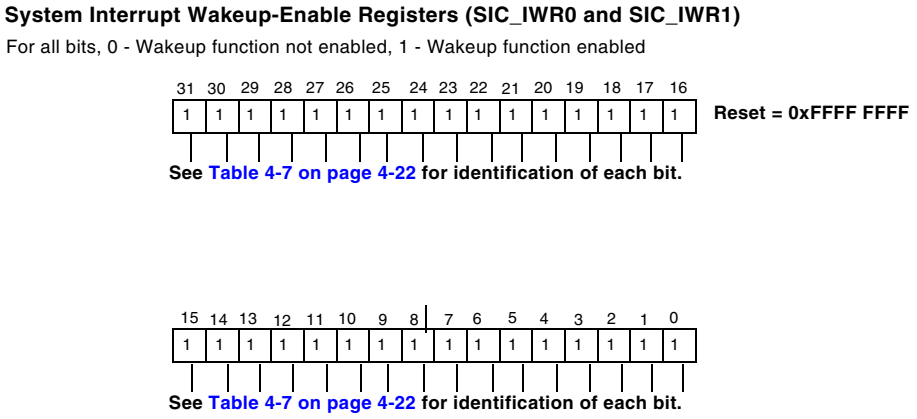

Figure 4-6. System Interrupt Wakeup-Enable Registers (SIC_IWR0 and SIC_IWR1)

System Interrupt Status Registers (SIC_ISRx)

The SIC includes two read-only System Interrupt Status registers, SIC_ISR0 and SIC_ISR1, shown in Figure 4-7. Each valid bit in each register corresponds to one of the peripheral interrupt sources. The bit is set when the SIC detects the interrupt is asserted and cleared when the SIC detects that the peripheral interrupt input has been deasserted. Note that for some peripherals, such as programmable flag asynchronous input interrupts, many cycles of latency may pass from the time that an interrupt service routine initiates the clearing of the interrupt (usually by writing a system MMR) to the time that the SIC senses that the interrupt has been deasserted.

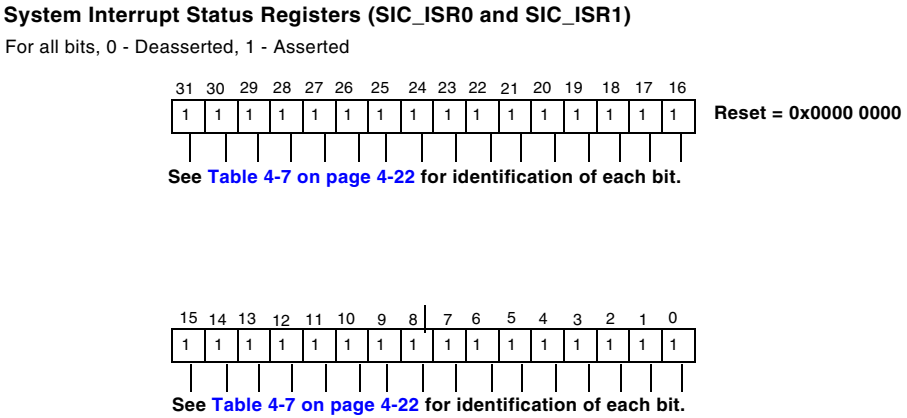

Figure 4-7. System Interrupt Status Registers (SIC_ISR0 and SIC_ISR1)

Depending on how interrupt sources map to the general-purpose interrupt inputs of the core, the interrupt service routine may have to interrogate multiple interrupt status bits to determine the source of the interrupt. One of the first instructions executed in an interrupt service routine should read `SIC_ISRx` to determine whether more than one of the peripherals sharing the input has asserted its interrupt output. The service routine should fully process all pending, shared interrupts before executing the RTI, which enables further interrupt generation on that interrupt input.

When an interrupt's service routine is finished, the `RTI` instruction clears the appropriate bit in the `IPEND` register. However, the relevant `SIC_ISRx` bit is not cleared unless the service routine clears the mechanism that generated the interrupt.

Many systems need relatively few interrupt-enabled peripherals, allowing each peripheral to map to a unique core priority level. In these designs, `SIC_ISRx` will seldom, if ever, need to be interrogated.

The SIC_ISRx registers are not affected by the state of the System Interrupt Mask registers (SIC_IMASKx) and can be read at any time. Writes to SIC_ISRx have no effect on their contents.

System Interrupt Mask Registers (SIC_IMASKx)

The System Interrupt Mask registers (SIC_IMASK0 and SIC_IMASK1) are shown in Figure 4-8. They allow masking of any peripheral interrupt source at the SIC, independently of whether it is enabled at the peripheral itself.

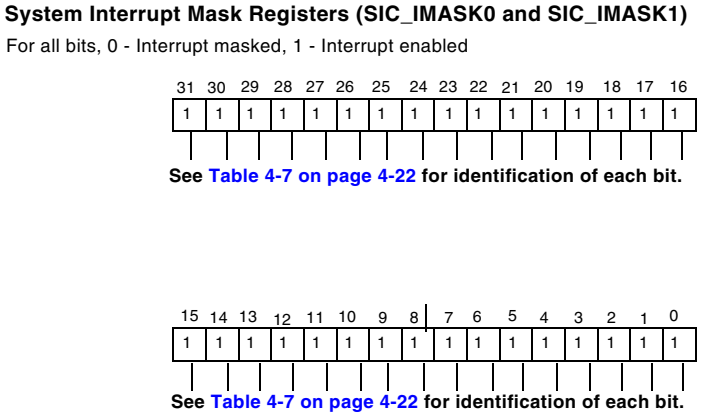

Figure 4-8. System Interrupt Mask Registers (SIC_IMASK0 and SIC_IMASK1)

A reset forces the contents of SIC_IMASKx to all 0s to mask off all peripheral interrupts. Writing a 1 to a bit location turns off the mask and enables the interrupt.

Although these registers can be read from or written to at any time (in Supervisor mode), they should be configured in the reset initialization sequence before enabling interrupts.

System Interrupt Assignment Registers (SIC_IARx)

The relative priority of peripheral interrupts can be set by mapping the peripheral interrupt to the appropriate general-purpose interrupt level in the core. The mapping is controlled by the System Interrupt Assignment register settings, as detailed in Figure 4-9 on page 4-35, Table 4-8, and Table 4-9. If more than one interrupt source is mapped to the same interrupt, they are logically OR'ed, with no hardware prioritization. Software can prioritize the interrupt processing as required for a particular system application.

Table 4-8. System Interrupt Assignment Register Reset Values

Interrupt	Reset Value
SIC_IAR0	0x0000 0000
SIC_IAR1	0x1111 1000
SIC_IAR2	0x2111 1111
SIC_IAR3	0x2222 2222
SIC_IAR4	0x3333 3222
SIC_IAR5	0x4333 3333
SIC_IAR6	0x2114 4444
SIC_IAR7	0x0000 6552

Table 4-9. Bit Position of Peripheral Interrupt Source in System Interrupt Controller Registers

Peripheral Interrupt Channel	Peripheral Interrupt Source	Bit Grouping of Peripheral Interrupt Source in System Interrupt Controller Registers							
		SIC_ IAR0	SIC_ IAR1	SIC_ IAR2	SIC_ IAR3	SIC_ IAR4	SIC_ IAR5	SIC_ IAR6	SIC_ IAR7
0	PLL Wakeup	3:0	-	-	-	-	-	-	-
1	DMA1 Error (generic)	7:4	-	-	-	-	-	-	-
2	DMA2 Error (generic)	11:8	-	-	-	-	-	-	-

Table 4-9. Bit Position of Peripheral Interrupt Source in System Interrupt Controller Registers (Cont'd)

Peripheral Interrupt Channel	Peripheral Interrupt Source	Bit Grouping of Peripheral Interrupt Source in System Interrupt Controller Registers							
		SIC_ IAR0	SIC_ IAR1	SIC_ IAR2	SIC_ IAR3	SIC_ IAR4	SIC_ IAR5	SIC_ IAR6	SIC_ IAR7
3	IMDMA Error	15:12	-	-	-	-	-	-	-
4	PPI0 Error	19:16	-	-	-	-	-	-	-
5	PPI1 Error	23:20	-	-	-	-	-	-	-
6	SPORT0 Error	27:24	-	-	-	-	-	-	-
7	SPORT1 Error	31:28	-	-	-	-	-	-	-
8	SPI Error	-	3:0	-	-	-	-	-	-
9	UART Error	-	7:4	-	-	-	-	-	-
10	Reserved	-	11:8	-	-	-	-	-	-
11	DMA1 Channel 0 interrupt (PPI0)	-	15:12	-	-	-	-	-	-
12	DMA1 Channel 1 interrupt (PPI1)	-	19:16	-	-	-	-	-	-
13	DMA1 Channel 2 interrupt	-	23:20	-	-	-	-	-	-
14	DMA1 Channel 3 interrupt	-	27:24	-	-	-	-	-	-
15	DMA1 Channel 4 interrupt	-	31:28	-	-	-	-	-	-
16	DMA1 Channel 5 interrupt	-	-	3:0	-	-	-	-	-
17	DMA1 Channel 6 interrupt	-	-	7:4	-	-	-	-	-
18	DMA1 Channel 7 interrupt	-	-	11:8	-	-	-	-	-
19	DMA1 Channel 8 interrupt	-	-	15:12	-	-	-	-	-
20	DMA1 Channel 9 interrupt	-	-	19:16	-	-	-	-	-
21	DMA1 Channel 10 interrupt	-	-	23:20	-	-	-	-	-
22	DMA1 Channel 11 interrupt	-	-	27:24	-	-	-	-	-
23	DMA2 Channel 0 interrupt (SPORT0 RX)	-	-	31:28	-	-	-	-	-

Table 4-9. Bit Position of Peripheral Interrupt Source in System Interrupt Controller Registers (Cont'd)

Peripheral Interrupt Channel	Peripheral Interrupt Source	Bit Grouping of Peripheral Interrupt Source in System Interrupt Controller Registers							
		SIC_ IAR0	SIC_ IAR1	SIC_ IAR2	SIC_ IAR3	SIC_ IAR4	SIC_ IAR5	SIC_ IAR6	SIC_ IAR7
24	DMA2 Channel 1 interrupt (SPORT0 TX)	-	-	-	3:0	-	-	-	-
25	DMA2 Channel 2 interrupt (SPORT1 RX)	-	-	-	7:4	-	-	-	-
26	DMA2 Channel 3 interrupt (SPORT1 TX)	-	-	-	11:8	-	-	-	-
27	DMA2 Channel 4 interrupt (SPI)	-	-	-	15:12	-	-	-	-
28	DMA2 Channel 5 interrupt (UART RX)	-	-	-	19:16	-	-	-	-
29	DMA2 Channel 6 interrupt (UART TX)	-	-	-	23:20	-	-	-	-
30	DMA2 Channel 7 interrupt	-	-	-	27:24	-	-	-	-
31	DMA2 Channel 8 interrupt	-	-	-	31:28	-	-	-	-
32	DMA2 Channel 9 interrupt	-	-	-	-	3:0	-	-	-
33	DMA2 Channel 10 interrupt	-	-	-	-	7:4	-	-	-
34	DMA2 Channel 11 interrupt	-	-	-	-	11:8	-	-	-
35	Timer0 interrupt	-	-	-	-	15:12	-	-	-
36	Timer1 interrupt	-	-	-	-	19:16	-	-	-
37	Timer2 interrupt	-	-	-	-	23:20	-	-	-
38	Timer3 interrupt	-	-	-	-	27:24	-	-	-
39	Timer4 interrupt	-	-	-	-	31:28	-	-	-
40	Timer5 interrupt	-	-	-	-	-	3:0	-	-
41	Timer6 interrupt	-	-	-	-	-	7:4	-	-
42	Timer7 interrupt	-	-	-	-	-	11:8	-	-

Table 4-9. Bit Position of Peripheral Interrupt Source in System Interrupt Controller Registers (Cont'd)

Peripheral Interrupt Channel	Peripheral Interrupt Source	Bit Grouping of Peripheral Interrupt Source in System Interrupt Controller Registers							
		SIC_ IAR0	SIC_ IAR1	SIC_ IAR2	SIC_ IAR3	SIC_ IAR4	SIC_ IAR5	SIC_ IAR6	SIC_ IAR7
43	Timer8 interrupt	-	-	-	-	-	15:12	-	-
44	Timer9 interrupt	-	-	-	-	-	19:16	-	-
45	Timer10 interrupt	-	-	-	-	-	23:20	-	-
46	Timer11 interrupt	-	-	-	-	-	27:24	-	-
47	Programmable Flags 15-0 interrupt A	-	-	-	-	-	31:28	-	-
48	Programmable Flags 15-0 interrupt B	-	-	-	-	-	-	3:0	-
49	Programmable Flags 31-16 interrupt A	-	-	-	-	-	-	7:4	-
50	Programmable Flags 31-16 interrupt B	-	-	-	-	-	-	11:8	-
51	Programmable Flags 47-32 interrupt A	-	-	-	-	-	-	15:12	-
52	Programmable Flags 47-32 interrupt B	-	-	-	-	-	-	19:16	-
53	DMA1 Channel 12/13 interrupt (Memory DMA Stream 0)	-	-	-	-	-	-	23:20	-
54	DMA1 Channel 14/15 interrupt (Memory DMA Stream 1)	-	-	-	-	-	-	27:24	-
55	DMA2 Channel 12/13 interrupt (Memory DMA Stream 0)	-	-	-	-	-	-	31:28	-
56	DMA2 Channel 14/15 interrupt (Memory DMA Stream 1)	-	-	-	-	-	-	-	3:0

Table 4-9. Bit Position of Peripheral Interrupt Source in System Interrupt Controller Registers (Cont'd)

Peripheral Interrupt Channel	Peripheral Interrupt Source	Bit Grouping of Peripheral Interrupt Source in System Interrupt Controller Registers							
		SIC_ IAR0	SIC_ IAR1	SIC_ IAR2	SIC_ IAR3	SIC_ IAR4	SIC_ IAR5	SIC_ IAR6	SIC_ IAR7
57	IMDMA Stream 0 interrupt	-	-	-	-	-	-	-	7:4
58	IMDMA Stream 1 interrupt	-	-	-	-	-	-	-	11:8
59	Watchdog Timer interrupt	-	-	-	-	-	-	-	15:12
60	Reserved	-	-	-	-	-	-	-	19:16
61	Reserved	-	-	-	-	-	-	-	23:20
62	Supplemental interrupt 0	-	-	-	-	-	-	-	27:24
63	Supplemental interrupt 1	-	-	-	-	-	-	-	31:28

For general-purpose interrupts with multiple peripheral interrupts assigned to them, take special care to ensure that software correctly processes all pending interrupts sharing that input. Software is responsible for prioritizing the shared interrupts.

These registers can be read or written at any time in Supervisor mode. It is advisable, however, to configure them in the Reset interrupt service routine before enabling interrupts. To prevent spurious or lost interrupt activity, these registers should be written to only when all peripheral interrupts are disabled.

System Interrupt Assignment Registers (SIC_IAR0 to SIC_IAR7)

Each 4-bit grouping represents the interrupt priority of a specific interrupt source (see Table 4-9 on page 4-30).

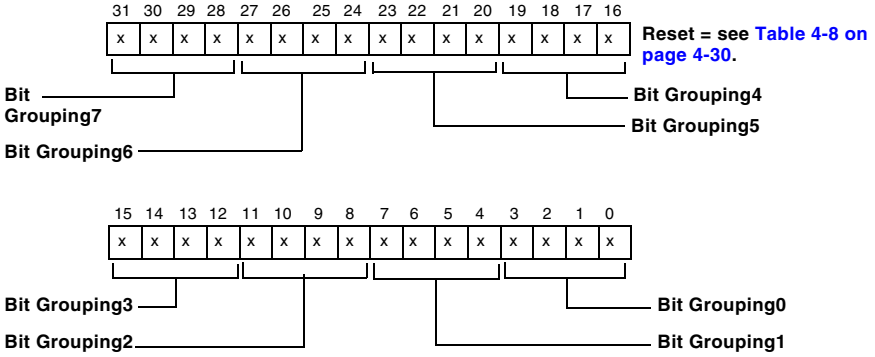

Figure 4-9. System Interrupt Assignment Registers (SIC_IAR0 to SIC_IAR7)

Table 4-10 defines the value to write in SIC_IARx to configure a peripheral for a particular IVG priority.

Table 4-10. IVG-Select Definition

General-Purpose Interrupt	Value in SIC_IARx
IVG7	0
IVG8	1
IVG9	2
IVG10	3
IVG11	4
IVG12	5
IVG13	6
IVG14	7
IVG15	8

Core Event Controller Registers

The Event Controller uses three 32-bit MMRs to coordinate pending event requests. In each of these MMRs, the lower 16 bits correspond to the 16 event levels (for example, bit 0 corresponds to “emulator mode”.

The registers are:

- IMASK - interrupt mask
- ILAT - interrupt latch
- IPEND - interrupt pending

These three registers are accessible in Supervisor Mode only.

Each core has its own separate IMASK, ILAT, and IPEND registers, but the MMR memory locations are aliased such that each core always accesses only its own set of these registers.

Core Interrupt Mask Register (IMASK)

This register indicates which interrupt levels are allowed to be serviced (see Figure 4-10). The IMASK register may be read and written in Supervisor mode. Bits15–5 have significance; bits4–0 are hard-coded to 1, and events of these levels are always enabled. If IMASK[N] == 1 and ILAT[N] == 1, then interrupt N will be taken if a higher priority is not already recognized. If IMASK[N] == 0, and ILAT[N] gets set by interrupt N, the interrupt will not be taken, and ILAT[N] will remain set.

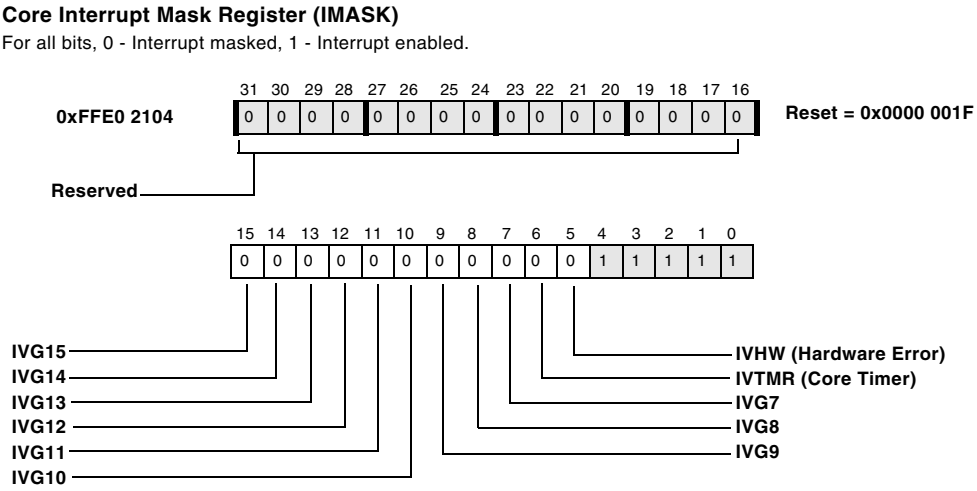

Figure 4-10. Core Interrupt Mask Register

Core Interrupt Latch Register (ILAT)

Each bit in `ILAT` indicates that the corresponding event is latched, but not yet accepted into the processor (see Figure 4-11). The bit is reset before the first instruction in the corresponding ISR is executed. At the point the interrupt is accepted, `ILAT[N]` will be cleared and `IPEND[N]` will be set simultaneously. The `ILAT` register can be read in Supervisor mode. Writes to `ILAT` are used to clear bits only (in Supervisor mode). To clear bit `N` from `ILAT`, first make sure that `IMASK[N] == 0`, and then write `ILAT[N] = 1`. This write functionality to `ILAT` is provided for cases where latched interrupt requests need to be cleared (cancelled) instead of serviced.

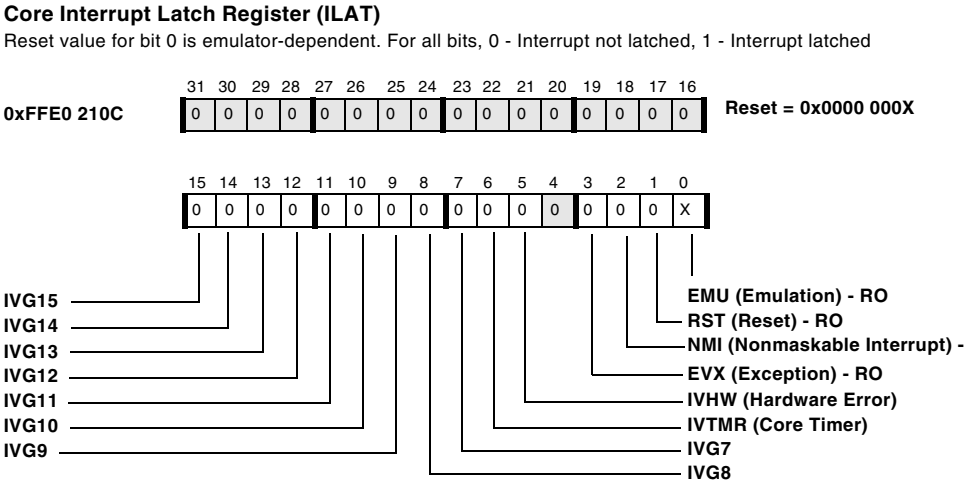

Figure 4-11. Core Interrupt Latch Register

The `RAISE` instruction can be used to set `ILAT[15]` through `ILAT[5]`, and also `ILAT[2]` or `ILAT[1]`.

Only the JTAG `TRST` pin can clear `ILAT[0]`.

Core Interrupt Pending Register (IPEND)

The `IPEND` register keeps track of all currently nested interrupts (see Figure 4-12). Each bit in `IPEND` indicates that the corresponding interrupt is currently active or nested at some level. It may be read in Supervisor mode, but not written. The `IPEND[4]` bit is used by the Event Controller to temporarily disable interrupts on entry and exit to an interrupt service routine.

When an event is processed, the corresponding bit in `IPEND` is set. The least significant bit in `IPEND` that is currently set indicates the interrupt that is currently being serviced. At any given time, `IPEND` holds the current status of all nested events.

Core Interrupt Pending Register (IPEND)
RO. For all bits except bit 4, 0 - No interrupt pending, 1 - Interrupt pending or active

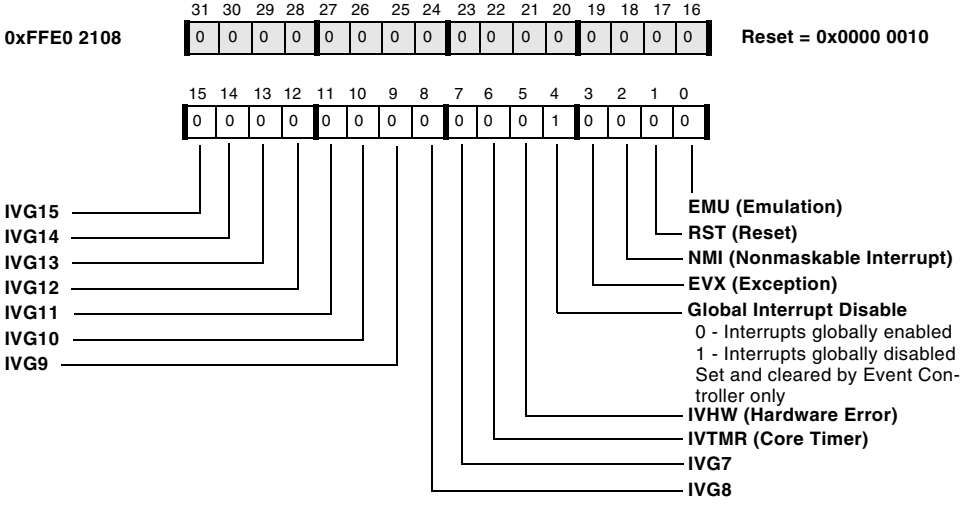

Figure 4-12. Core Interrupt Pending Register

Global Enabling/Disabling of Interrupts

General-purpose interrupts can be globally disabled with the `CLI Dreg` instruction and re-enabled with the `STI Dreg` instruction, both of which are only available in Supervisor mode. Reset, NMI, emulation, and exception events cannot be globally disabled. Globally disabling interrupts clears `IMASK[15:5]` after saving `IMASK`'s current state. See “Enable Interrupts” and “Disable Interrupts” in the “External Event Management” chapter of the *Blackfin Processor Programming Reference*.

When program code is too time-critical to be delayed by an interrupt, disable general-purpose interrupts, but be sure to re-enable them at the conclusion of the code sequence.

Event Vector Table

The Event Vector Table (EVT) is a hardware table with sixteen entries that are each 32 bits wide. The EVT contains an entry for each possible core event. Entries are accessed as MMRs, and each entry can be programmed at reset with the corresponding vector address for the interrupt service routine. When an event occurs, instruction fetch starts at the address location in the EVT entry for that event.

Each core has its own separate Event Vector Table MMRs (EVT15–EVT0) but the MMR memory locations are aliased such that each core always accesses only its own set of these registers.

The processor architecture allows unique addresses to be programmed into each of the interrupt vectors; that is, interrupt vectors are not determined by a fixed offset from an interrupt vector table base address. This approach minimizes latency by not requiring a long jump from the vector table to the actual ISR code.

Table 4-11 lists events by priority. Each event has a corresponding bit in the event state registers `ILAT`, `IMASK`, and `IPEND`.

Table 4-11. Core Event Vector Table

Event Number	Event Class	Name	MMR Location	Notes
EVT0	Emulation	EMU	0xFFE0 2000	Highest priority. Vector address is provided by JTAG.
EVT1	Reset	RST	0xFFE0 2004	Read-only.
EVT2	NMI	NMI	0xFFE0 2008	
EVT3	Exception	EVX	0xFFE0 200C	
EVT4	Reserved	Reserved	0xFFE0 2010	Reserved vector.
EVT5	Hardware Error	IVHW	0xFFE0 2014	
EVT6	Core Timer	IVTMR	0xFFE0 2018	
EVT7	Interrupt 7	IVG7	0xFFE0 201C	

Table 4-11. Core Event Vector Table (Cont'd)

Event Number	Event Class	Name	MMR Location	Notes
EVT8	Interrupt 8	IVG8	0xFFE0 2020	
EVT9	Interrupt 9	IVG9	0xFFE0 2024	
EVT10	Interrupt 10	IVG10	0xFFE0 2028	
EVT11	Interrupt 11	IVG11	0xFFE0 202C	
EVT12	Interrupt 12	IVG12	0xFFE0 2030	
EVT13	Interrupt 13	IVG13	0xFFE0 2034	
EVT14	Interrupt 14	IVG14	0xFFE0 2038	
EVT15	Interrupt 15	IVG15	0xFFE0 203C	Lowest priority.

Emulation

An emulation event causes the processor to enter Emulation mode, in which instructions are read from the JTAG interface. It is the highest-priority interrupt to the core.

For detailed information about emulation, see Chapter 18, "Blackfin Processor Debug".

Reset

The reset interrupt (RST) can be initiated via the RESET pin or through expiration of the Watchdog timer. This location differs from that of other interrupts in that its content is read-only. Writes to this address change the register but do not change where the processor vectors upon reset. For more information, see "Reset State" on page 3-10 and "Booting Methods" on page 3-19.

The core has an output that indicates that a double-fault has occurred. This is a non-recoverable state. The system (via the SWRST register) can be programmed to send a reset request if a double-fault condition is detected.

Subsequently, the reset request forces a system reset for core and peripherals.

The reset vector is determined by the processor system. It points to the start of the on-chip boot ROM, or to the start of external asynchronous memory, depending on the state of the BMODE[1:0] pins. Refer to Table 4-12.

Table 4-12. Reset Vector Addresses

Boot Source	BMODE[1:0]	Execution Start Address
Bypass boot ROM; execute from 16-bit-wide external memory (Async Bank 0)	00	0x2000 0000
Use boot ROM to boot from 8/16-bit flash	01	0xEF00 0000
Use boot ROM to boot from SPI host	10	0xEF00 0000
Use boot ROM to configure and load boot code from SPI serial EEPROM (16-, or 24-bit address range)	11	0xEF00 0000

If the BMODE[1:0] pins indicate either booting from flash or serial EEPROM, the reset vector points to the start of the internal boot ROM, where a small bootstrap kernel resides. The bootstrap code reads the System Reset Configuration register (SYSCR) to determine the value of the BMODE[1:0] pins, which determine the appropriate boot sequence. For information about the ADSP-BF561 boot ROM, see "Booting Methods" on page 3-19.

If the BMODE[1:0] pins indicate to bypass boot ROM, the reset vector points to the start of the external asynchronous memory region. In this mode, the internal boot ROM is not used. To support reads from this memory region, the external bus interface unit (EBIU) uses the default external memory configuration that results from hardware reset.

NMI (Non-Maskable Interrupt)

The NMI entry is reserved for a non-maskable interrupt, which can be generated by the Watchdog timer or by the NMI input signal to the processor. An example of an event that requires immediate processor attention, and thus is appropriate as an NMI, is a power-down warning.

If an exception occurs in an event handler that is already servicing an exception, NMI, reset, or emulation event, this will trigger a double fault condition, and the address of the excepting instruction will be written to RETX.

Exceptions

Exceptions are synchronous to the instruction stream. In other words, a particular instruction causes an exception when it attempts to finish execution. No instructions after the offending instruction are executed before the exception handler takes effect.

Many of the exceptions are memory related. For example, an exception is given when a misaligned access is attempted, or when a cacheability protection lookaside buffer (CPLB) miss or protection violation occurs. Exceptions are also given when illegal instructions or illegal combinations of registers are executed.

An excepting instruction may or may not commit before the exception event is taken, depending on if it is a service type or an error type exception.

An instruction causing a service type event will commit, and the address written to the RETX register will be the next instruction after the excepting one. An example of a service type exception is the single step.

An instruction causing an error type event cannot commit, so the address written to the RETX register will be the address of the offending instruction. An example of an error type event is a CPLB miss.

Usually the RETX register contains the correct address to return to. To skip over an excepting instruction, take care in case the next address is not simply the next linear address. This could happen when the excepting instruction is a loop end. In that case, the proper next address would be the loop top.

The EXCAUSE[5:0] field in the Sequencer Status register (SEQSTAT) is written whenever an exception is taken, and indicates to the exception handler which type of exception occurred. Refer to Table 4-13 for a list of events that cause exceptions.

If an exception occurs in an event handler that is already servicing an exception, NMI, reset, or emulation event, this will trigger a double fault condition, and the address of the excepting instruction will be written to RETX.

Table 4-13. Events That Cause Exceptions

Exception	EXCAUSE[5:0]	Type: (E) Error (S) Service See table note.	Notes/Examples
Force Exception instruction EXCPT with 4-bit m field	m field	S	Instruction provides 4 bits of EXCAUSE.
Single step	0x10	S	When the processor is in single step mode, every instruction generates an exception. Primarily used for debugging.
Note: For services (S), the return address is the address of the instruction that follows the exception. For errors (E), the return address is the address of the excepting instruction.			

Table 4-13. Events That Cause Exceptions (Cont'd)

Exception	EXCAUSE[5:0]	Type: (E) Error (S) Service See table note.	Notes/Examples
Exception caused by a trace buffer full condition	0x11	S	The processor takes this exception when the trace buffer overflows (only when enabled by the Trace Unit Control register).
Undefined instruction	0x21	E	May be used to emulate instructions that are not defined for a particular processor implementation.
Illegal instruction combination	0x22	E	See section for multi-issue rules in the *Blackfin Processor Programming Reference*.
Data access CPLB protection violation	0x23	E	Attempted read or write to Supervisor resource, or illegal data memory access. Supervisor resources are registers and instructions that are reserved for Supervisor use: Supervisor-only registers, all MMRs, and Supervisor-only instructions. (A simultaneous, dual access to two MMRs using the data address generators generates this type of exception.) In addition, this entry is used to signal a protection violation caused by disallowed memory access, and it is defined by the Memory Management Unit (MMU) cacheability protection lookaside buffer (CPLB). Also caused by a data access to a reserved L1 or L1 instruction memory.
Data access misaligned address violation	0x24	E	Attempted misaligned data memory or data cache access.
Unrecoverable event	0x25	E	For example, an exception generated while processing a previous exception.

Note: For services (S), the return address is the address of the instruction that follows the exception. For errors (E), the return address is the address of the excepting instruction.

Table 4-13. Events That Cause Exceptions (Cont'd)

Exception	EXCAUSE[5:0]	Type: (E) Error (S) Service See table note.	Notes/Examples
Data access CPLB miss	0x26	E	Used by the MMU to signal a CPLB miss on a data access.
Data access multiple CPLB hits	0x27	E	More than one CPLB entry matches data fetch address.
Exception caused by an emulation watchpoint match	0x28	E	There is a watchpoint match, and one of the EMUSW bits in the Watchpoint Instruction Address Control register (WPIACTL) is set.
Instruction fetch misaligned address violation	0x2A	E	Attempted misaligned instruction cache fetch. On a misaligned instruction fetch exception, the return address provided in RETX is the destination address which is misaligned, rather than the address of the offending instruction. For example, if an indirect branch to a misaligned address held in P0 is attempted, the return address in RETX is equal to P0, rather than to the address of the branch instruction. (Note this exception can never be generated from PC-relative branches, only from indirect branches.)
Instruction fetch CPLB protection violation	0x2B	E	Illegal instruction fetch access (memory protection violation). Also caused by an instruction fetch from reserved L1 or L1 data memory.
Instruction fetch CPLB miss	0x2C	E	CPLB miss on an instruction fetch.

Note: For services (S), the return address is the address of the instruction that follows the exception. For errors (E), the return address is the address of the excepting instruction.

Table 4-13. Events That Cause Exceptions (Cont'd)

Exception	EXCAUSE[5:0]	Type: (E) Error (S) Service See table note.	Notes/Examples
Instruction fetch multiple CPLB hits	0x2D	E	More than one CPLB entry matches instruction fetch address.
Illegal use of supervisor resource	0x2E	E	Attempted to use a Supervisor register or instruction from User mode. Supervisor resources are registers and instructions that are reserved for Supervisor use: Supervisor-only registers, all MMRs, and Supervisor-only instructions.
Note: For services (S), the return address is the address of the instruction that follows the exception. For errors (E), the return address is the address of the excepting instruction.			

If an instruction causes multiple exceptions, only the exception with the highest priority is taken. Table 4-14 ranks exceptions by descending priority.

Table 4-14. Exceptions by Descending Priority

Priority	Exception	EXCAUSE
1	Unrecoverable Event	0x25
2	I-Fetch Multiple CPLB Hits	0x2D
3	I-Fetch Misaligned Access	0x2A
4	I-Fetch Protection Violation	0x2B
5	I-Fetch CPLB Miss	0x2C
6	I-Fetch Access Exception	0x29
7	Watchpoint Match	0x28
8	Undefined Instruction	0x21
9	Illegal Combination	0x22
10	Illegal Use of Protected Resource	0x2E

Table 4-14. Exceptions by Descending Priority (Cont'd)

Priority	Exception	EXCAUSE
11	DAG0 Multiple CPLB Hits	0x27
12	DAG0 Misaligned Access	0x24
13	DAG0 Protection Violation	0x23
14	DAG0 CPLB Miss	0x26
15	DAG1 Multiple CPLB Hits	0x27
16	DAG1 Misaligned Access	0x24
17	DAG1 Protection Violation	0x23
18	DAG1 CPLB Miss	0x26
19	EXCPT Instruction	m field
20	Single Step	0x10
21	Trace Buffer	0x11

Exceptions While Executing an Exception Handler

While executing the exception handler, avoid issuing an instruction that generates another exception. If an exception is caused while executing code within the exception handler, the NMI handler, the reset vector, or in emulator mode:

- The excepting instruction is not committed. All writebacks from the instruction are prevented.
- The generated exception is not taken.
- The EXCAUSE field in SEQSTAT is updated with an unrecoverable event code.
- The address of the offending instruction is saved in RETX. Note if the processor were executing, for example, the NMI handler, the RETN register would not have been updated; the excepting instruction address is always stored in RETX.

To determine whether an exception occurred while an exception handler was executing, check SEQSTAT at the end of the exception handler for the code indicating an "unrecoverable event" (EXCAUSE = 0x25). If an unrecoverable event occurred, register RETX holds the address of the most recent instruction to cause an exception. This mechanism is not intended for recovery, but rather for detection.

Hardware Error Interrupt

The Hardware Error Interrupt indicates a hardware error or system malfunction. Hardware errors occur when logic external to the core, such as a memory bus controller, is unable to complete a data transfer (read or write) and asserts the core's error input signal. Such hardware errors invoke the Hardware Error Interrupt (interrupt IVHW in the Event Vector Table (EVT) and ILAT, IMASK, and IPEND registers). The Hardware Error Interrupt service routine can then read the cause of the error from the 5-bit HWERRCAUSE field appearing in the Sequencer Status register (SEQSTAT) and respond accordingly.

The Hardware Error Interrupt is generated by:

- Bus parity errors
- Internal error conditions within the core, such as Performance Monitor overflow
- Peripheral errors
- Bus timeout errors

The list of supported hardware conditions, with their related HWERRCAUSE codes, appears in Table 4-15. The bit code for the most recent error appears in the HWERRCAUSE field. If multiple hardware errors occur simultaneously, only the last one can be recognized and serviced. The core does not support prioritizing, pipelining, or queuing multiple error codes. The

Hardware Error Interrupt remains active as long as any of the error conditions remain active.

Table 4-15. Conditions Causing Hardware Error Interrupts

Hardware Condition	HWERRCAUSE (binary)	HWERRCAUSE (hexadecimal)	Notes / Examples
System MMR Error	0b00010	0x02	An error can occur if an invalid System MMR location is accessed, if a 32-bit register is accessed with a 16-bit instruction, or if a 16-bit register is accessed with a 32-bit instruction.
External Memory Addressing Error	0b00011	0x03	
Performance Monitor Overflow	0b10010	0x12	
RAISE 5 instruction	0b11000	0x018	Software issued a RAISE 5 instruction to invoke the Hardware Error Interrupt (IVHW).
Reserved	All other bit combinations.	All other values.	

Core Timer

The Core Timer Interrupt (`IVTMR`) is triggered when the core timer value reaches zero. See Chapter 15, “Timers”.

General-Purpose Interrupts (IVG7-IVG15)

General-purpose interrupts are used for any event that requires processor attention. For instance, a DMA controller may use them to signal the end of a data transmission, or a serial communications device may use them to signal transmission errors.

Software can also trigger general-purpose interrupts by using the RAISE instruction. The RAISE instruction forces events for interrupts IVG15-IVG7, IVTMR, IVHW, NMI, and RST, but not for exceptions and emulation (EVX and EMU, respectively).

It is recommended to reserve the two lowest priority interrupts (IVG15 and IVG14) for software interrupt handlers.

Servicing Interrupts

The Core Event Controller (CEC) has a single interrupt queueing element per event—a bit in the ILAT register. The appropriate ILAT bit is set when an interrupt rising edge is detected (which takes two core clock cycles) and cleared when the respective IPEND register bit is set. The IPEND bit indicates that the event vector has entered the core pipeline. At this point, the CEC recognizes and queues the next rising edge event on the corresponding interrupt input. The minimum latency from the rising edge transition of the general-purpose interrupt to the IPEND output assertion is three core clock cycles. However, the latency can be much higher, depending on the core's activity level and state.

To determine when to service an interrupt, the controller logically ANDs the three quantities in ILAT, IMASK, and the current processor priority level.

Servicing the highest priority interrupt involves these actions:

1. The interrupt vector in the Event Vector Table (EVT) becomes the next fetch address.

   On an interrupt, most instructions currently in the pipeline are aborted. On a service exception, all instructions after the excepting instruction are aborted. On an error exception, the excepting instruction and all instructions after it are aborted.

2. The return address is saved in the appropriate return register.

   The return register is RETI for interrupts, RETX for exceptions, RETN for NMIs, and RETE for debug emulation. The return address is the address of the instruction after the last instruction executed from normal program flow.

3. Processor mode is set to the level of the event taken.

   If the event is an NMI, exception, or interrupt, the processor mode is Supervisor. If the event is an emulation exception, the processor mode is Emulation.

4. Before the first instruction starts execution, the corresponding interrupt bit in ILAT is cleared and the corresponding bit in IPEND is set.

   Bit IPEND[4] is also set to disable all interrupts until the return address in RETI is saved.

Nesting of Interrupts

Interrupts are handled either with or without nesting.

Non-Nested Interrupts

If interrupts do not require nesting, all interrupts are disabled during the interrupt service routine (ISR). Note, however, that emulation, NMI, and exceptions are still accepted by the system.

When the system does not need to support nested interrupts, there is no need to store the return address held in RETI. Only the portion of the machine state used in the interrupt service routine must be saved in the Supervisor stack. To return from a non-nested interrupt service routine,

only the RTI instruction must be executed, because the return address is already held in the RETI register.

Figure 4-13 shows an example of interrupt handling where interrupts are globally disabled for the entire interrupt service routine.

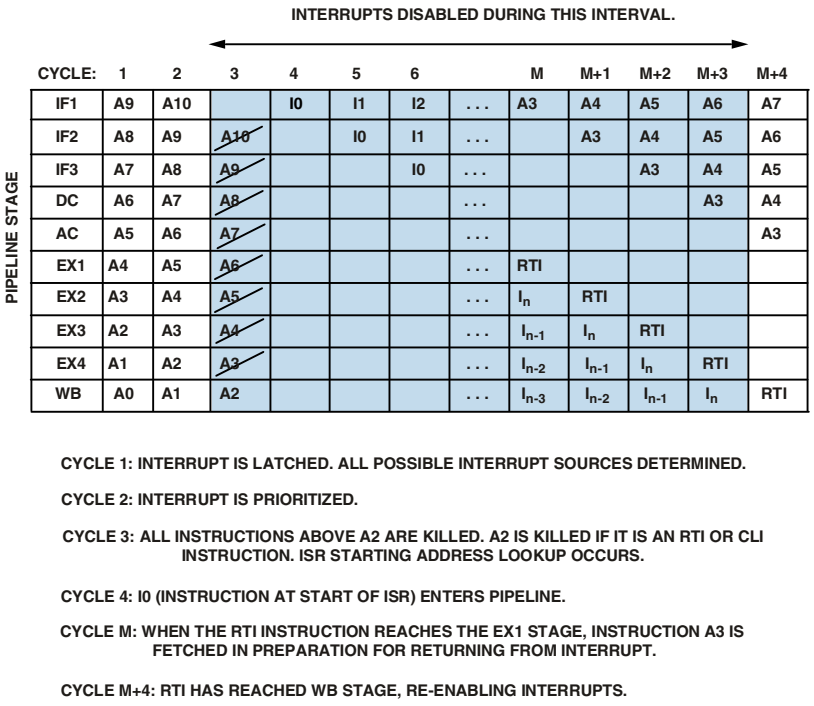

Figure 4-13. Non-Nested Interrupt Handling

Nested Interrupts

If interrupts require nesting, the return address to the interrupted point in the original ISR must be explicitly saved and subsequently restored when execution of the nested ISR has completed. The first instruction in an interrupt service routine that supports nesting must save the return address currently held in RETI by pushing it onto the Supervisor stack

([--SP]=RETI). This clears the global interrupt disable bit IPEND[4], enabling interrupts. Next, all registers that are modified by the interrupt service routine are saved onto the Supervisor stack. Processor state is stored in the Supervisor stack, not in the User stack. Hence, the instructions to push RETI ([--SP]=RETI) and pop RETI (RETI=[SP++]) use the Supervisor stack.

Figure 4-14 illustrates that by pushing RETI onto the stack, interrupts can be re-enabled during an ISR, resulting in a short duration where interrupts are globally disabled.

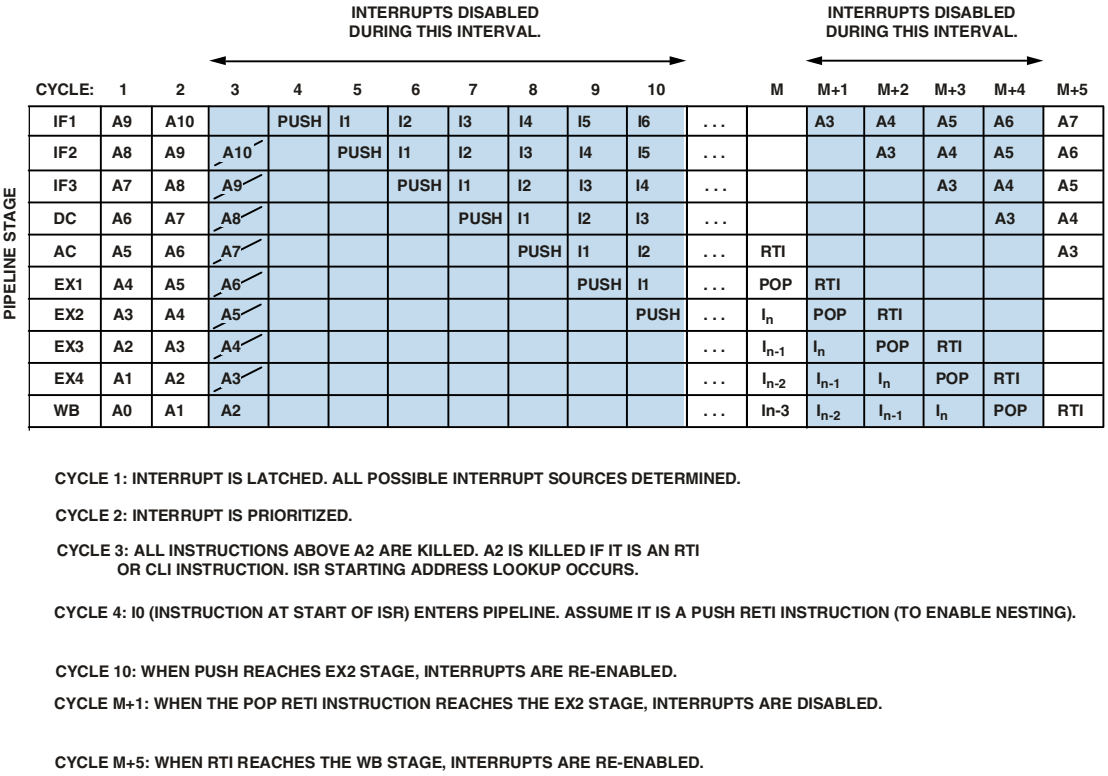

Figure 4-14. Nested Interrupt Handling

Example Prolog Code for Nested Interrupt Service Routine

Listing 4-2. Prolog Code for Nested ISR

```
/* Prolog code for nested interrupt service routine.
Push return address in RETI into Supervisor stack, ensuring that
interrupts are back on. Until now, interrupts have been
suspended.*/
ISR:
[--SP] = RETI; /* Enables interrupts and saves return address to
stack */
[--SP] = ASTAT;
[--SP] = FP;
[-- SP] = (R7:0, P5:0);
/* Body of service routine. Note that none of the processor
resources (accumulators, DAGs, loop counters and bounds) have
been saved. It's assumed that this interrupt service routine does
not use them. */
```

Example Epilog Code for Nested Interrupt Service Routine

Listing 4-3. Epilog Code for Nested ISR

```
/* Epilog code for nested-interrupt service routine.
Restore ASTAT, Data and Pointer registers. Popping RETI from
Supervisor stack ensures that interrupts are suspended between
load of return address and RTI. */
(R7:0, P5:0) = [SP++];
FP    = [SP++];
ASTAT = [SP++];
RETI = [SP++];
```

```
/* Execute RTI, which jumps to return address, re-enables inter-
rupts, and switches to User mode if this is the last nested
interrupt in service. */
RTI;
```

The RTI instruction causes the return from an interrupt. The return address is popped into the RETI register from the stack, an action that suspends interrupts from the time that RETI is restored until RTI finishes executing. The suspension of interrupts prevents a subsequent interrupt from corrupting the RETI register.

Next, the RTI instruction clears the highest-priority bit that is currently set in IPEND. The processor then jumps to the address pointed to by the value in the RETI register and re-enables interrupts by clearing IPEND[4].

Logging of Nested Interrupt Requests

The System Interrupt Controller (SIC) detects level-sensitive interrupt requests from the peripherals. The Core Event Controller (CEC) provides edge-sensitive detection for its general-purpose interrupts (IVG7-IVG15). Consequently, the SIC generates a synchronous interrupt pulse to the CEC and then waits for interrupt acknowledgement from the CEC. When the interrupt has been acknowledged by the core (via assertion of the appropriate IPEND output), the SIC generates another synchronous interrupt pulse to the CEC if the peripheral interrupt is still asserted. This way, the system does not lose peripheral interrupt requests that occur during servicing of another interrupt.

Multiple interrupt sources can map to a single core-processor general-purpose interrupt. Because of this, multiple pulse assertions from the SIC can occur simultaneously, before, or during interrupt processing for an interrupt event that is already detected on this interrupt input. For a shared interrupt, the IPEND interrupt-acknowledge mechanism described above re-enables all shared interrupts. If any of the shared interrupt sources are still asserted, at least one pulse is again generated by the SIC. The

Interrupt Status registers indicate the current state of the shared interrupt sources.

Exception Handling

Interrupts and exceptions treat instructions in the pipeline differently:

- When an interrupt occurs, all instructions in the pipeline are aborted.
- When an exception occurs, all instructions in the pipeline after the excepting instruction are aborted. For service exceptions, the excepting instruction is also aborted.

Because exceptions, NMIs, and emulation events have a dedicated return register, guarding the return address is optional. Consequently, the `PUSH` and `POP` instructions for exceptions, NMIs, and emulation events do not affect the interrupt system.

Note, however, that the return instructions for exceptions (`RTX`, `RTN`, and `RTE`) do clear the least significant bit currently set in `IPEND`.

Deferring Exception Processing

Exception handlers are usually long routines, because they must discriminate among several exception causes and take corrective action accordingly. The length of the routines may result in long periods during which the interrupt system is, in effect, suspended.

To avoid lengthy suspension of interrupts, write the exception handler to identify the exception cause, but defer the processing to a low-priority interrupt. To set up the low-priority interrupt handler, use the Force Interrupt / Reset instruction (`RAISE`).

When deferring the processing of an exception to lower-priority interrupt IVGx, the system must guarantee that IVGx is entered before returning to the application-level code that issued the exception. If a pending interrupt of higher priority than IVGx occurs, it is acceptable to enter the high-priority interrupt before IVGx.

Example Code for an Exception Handler

Listing 4-4 is for an exception routine handler with deferred processing.

Listing 4-4. Exception Routine Handler With Deferred Processing

```
/* Determine exception cause by examining EXCAUSE field in
SEQSTAT (first save contents of R0, P0, P1 and ASTAT in Supervi-
sor SP) */
[--SP] = R0;
[--SP] = P0;
[--SP] = P1;
[--SP] = ASTAT;
R0 = SEQSTAT;
/* Mask the contents of SEQSTAT, and leave only EXCAUSE in R0 */
R0 <<= 26;
R0 >>= 26;
/* Using jump table EVTABLE, jump to the event pointed to by R0
*/
P0 = R0;
P1 = _EVTABLE;
P0 = P1 + (P0 << 1);
R0 = W [P0] (Z);
P1 = R0;
JUMP (PC + P1);
/* The entry point for an event is as follows. Here, processing
is deferred to low-priority interrupt IVG15. Also, parame-
ter-passing would typically be done here. */
```

```
_EVENT1:
RAISE 15;
JUMP.S _EXIT;
/* Entry for event at IVG14 */
_EVENT2:
RAISE 14;
JUMP.S _EXIT;
/* comments for other events */
/* At the end of handler, restore R0, P0, P1 and ASTAT, and
return. */
_EXIT:
ASTAT = [SP++];
P1 = [SP++];
P0 = [SP++];
R0 = [SP++];
RTX;
_EVTABLE:
.byte2 addr_event1;
.byte2 addr_event2;
...
.byte2 addr_eventN;
/* The jump table EVTABLE holds 16-bit address offsets for each
event. With offsets, this code is position-independent and the
table is small.
+--------------+
| addr_event1 | _EVTABLE
+--------------+
| addr_event2 | _EVTABLE + 2
+--------------+
|      …       |
+--------------+
| addr_eventN | _EVTABLE + 2N
+--------------+
*/
```

Example Code for an Exception Routine

Listing 4-5 provides an example framework for an interrupt routine jumped to from an exception handler such as that described above.

Listing 4-5. Interrupt Routine for Handling Exception

```
[--SP] = RETI;   /* Push return address on stack. */

/* Put body of routine here.*/

RETI = [SP++];   /* To return, pop return address and jump.*/

RTI;   /* Return from interrupt. */
```

Example Code for Using Hardware Loops in an ISR

Listing 4-6 shows the optimal method of saving and restoring when using hardware loops in an interrupt service routine.

Listing 4-6. Saving and Restoring Hardware Loops

```
lhandler:
<save other registers here>
[--SP] = LC0; /* save loop 0 */
[--SP] = LB0;
[--SP] = LT0;

<handler code here>
```

```
/* If the handler uses loop 0, it's a good idea to have
it leave LC0 equal to zero at the end. Normally, this will
happen naturally as a loop is fully executed. If LC0 == 0,
then LT0 and LB0 restores will not incur additional cycles.
If LC0!= 0 when the following pops happen, each pop will
incur a 10-cycle "replay" penalty. Popping or writing LC0
always incurs the penalty. */

LT0 = [SP++];
LB0 = [SP++];
LC0 = [SP++];   /* this will cause a "replay," that is, a
                   10-cycle refetch. */

<restore other registers here>

RTI;
```

Additional Usability Issues

The following sections describe additional usability issues.

Executing RTX, RTN, or RTE in a Lower Priority Event

Instructions `RTX`, `RTN`, and `RTE` are designed to return from an exception, NMI, or emulator event, respectively. Do not use them to return from a lower priority event. To return from an interrupt, use the `RTI` instruction. Failure to use the correct instruction may produce unintended results.

In the case of `RTX`, bit `IPEND[3]` is cleared. In case of `RTI`, the bit of the highest priority interrupt in `IPEND` is cleared.

Recommendation for Allocating the System Stack

The software stack model for processing exceptions implies that the Supervisor stack must never generate an exception while the exception handler is saving its state. However, if the Supervisor stack grows past a CPLB entry or SRAM block, it may, in fact, generate an exception.

To guarantee that the Supervisor stack never generates an exception—never overflows past a CPLB entry or SRAM block while executing the exception handler—calculate the maximum space that all interrupt service routines and the exception handler occupy while they are active, and then allocate this amount of SRAM memory.

Latency in Servicing Events

In some processor architectures, if instructions are executed from external memory and an interrupt occurs while the instruction fetch operation is underway, then the interrupt is held off from being serviced until the current fetch operation has completed. Consider a processor operating at 300MHz and executing code from external memory with 100ns access times. Depending on when the interrupt occurs in the instruction fetch operation, the interrupt service routine (ISR) may be held off for around 30 instruction clock cycles. When cache line fill operations are taken into account, the ISR could be held off for many hundreds of cycles.

In order for high-priority interrupts to be serviced with the least latency possible, the processor allows any high-latency fill operation to be completed at the system level, while an ISR executes from L1 memory. Figure 4-15 illustrates this concept.

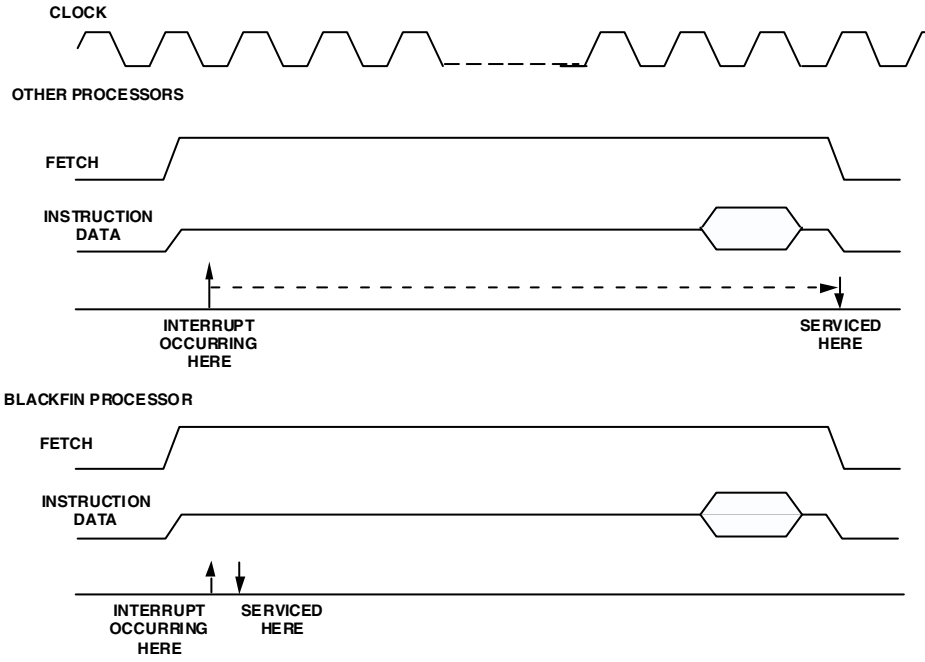

Figure 4-15. Minimizing Latency in Servicing an ISR

If an instruction load operation misses the L1 Instruction Cache and generates a high-latency line fill operation, then when an interrupt occurs, it is not held off until the fill has completed. Instead, the processor executes the ISR in its new context, and the cache fill operation completes in the background.

Note that the ISR must reside in L1 cache or SRAM memory and must not generate a cache miss, an L2 memory access, or a peripheral access, as the processor is already busy completing the original cache line fill operation. If a load or store operation is executed in the ISR requiring one of these accesses, then the ISR is held off while the original external access is completed, before initiating the new load or store.

If the ISR finishes execution before the load operation has completed, then the processor continues to stall, waiting for the fill to complete.

This same behavior is also exhibited for stalls involving reads of slow data memory or peripherals.

Writes to slow memory generally do not show this behavior, as the writes are deemed to be single-cycle, being immediately transferred to the write buffer for subsequent execution.

For detailed information about cache and memory structures, see Chapter 6, “Memory”.

5 DATA ADDRESS GENERATORS

The Data Address Generators (DAGs) generate addresses for data moves to and from memory. By generating addresses, the DAGs let programs refer to addresses indirectly, using a DAG register instead of an absolute address.

The DAG architecture, shown in Figure 5-1, supports several functions that minimize overhead in data access routines. These functions include:

- Supply address – Provides an address during a data access
- Supply address and post-modify – Provides an address during a data move and auto-increments/decrements the stored address for the next move
- Supply address with offset – Provides an address from a base with an offset without incrementing the original address pointer
- Modify address – Increments or decrements the stored address without performing a data move
- Bit-reversed carry address – Provides a bit-reversed carry address during a data move without reversing the stored address

The DAG subsystem comprises two DAG Arithmetic units, nine Pointer registers, four Index registers and four complete sets of related Modify, Base, and Length registers. These registers hold the values that the DAGs use to generate addresses. The types of registers are:

- Index registers, `I[3:0]`. Unsigned 32-bit Index registers hold an address pointer to memory. For example, the instruction `R3 = [I0]` loads the data value found at the memory location pointed to by the register `I0`. Index registers can be used for 16- and 32-bit memory accesses.
- Modify registers, `M[3:0]`. Signed 32-bit Modify registers provide the increment or step size by which an Index register is post-modified during a register move.

  For example, the `R0 = [I0 ++ M1]` instruction directs the DAG to:

  – Output the address in register `I0`

  – Load the contents of the memory location pointed to by `I0` into `R0`

  – Modify the contents of `I0` by the value contained in the `M1` register
- Base and Length registers, `B[3:0]` and `L[3:0]`. Unsigned 32-bit Base and Length registers set up the range of addresses and the starting address of a circular buffer. Each `B`, `L` pair is always coupled with a corresponding I-register, for example, `I3`, `B3`, `L3`. For more information on circular buffers, see "Addressing Circular Buffers" on page 5-6.
- Pointer registers, `P[5:0]`, `FP`, `USP`, and `SP`. 32-bit Pointer registers hold an address pointer to memory. The `P[5:0]` field, `FP` (Frame Pointer) and `SP/USP` (Stack Pointer/User Stack Pointer) can be manipulated and used in various instructions. For example, the instruction `R3 = [P0]` loads the register `R3` with the data value

found at the memory location pointed to by the register P0. The Pointer registers have no effect on circular buffer addressing. They can be used for 8-, 16-, and 32-bit memory accesses. For added mode protection, SP is accessible only in Supervisor mode, while USP is accessible in User mode.

Do not assume the L-registers are automatically initialized to zero for linear addressing. The I-, M-, L-, and B-registers contain random values after reset. For each I-register used, programs must initialize the corresponding L-registers to zero for linear addressing or to the buffer length for circular buffer addressing.

Note all DAG registers must be initialized individually. Initializing a B-register does not automatically initialize the I-register.

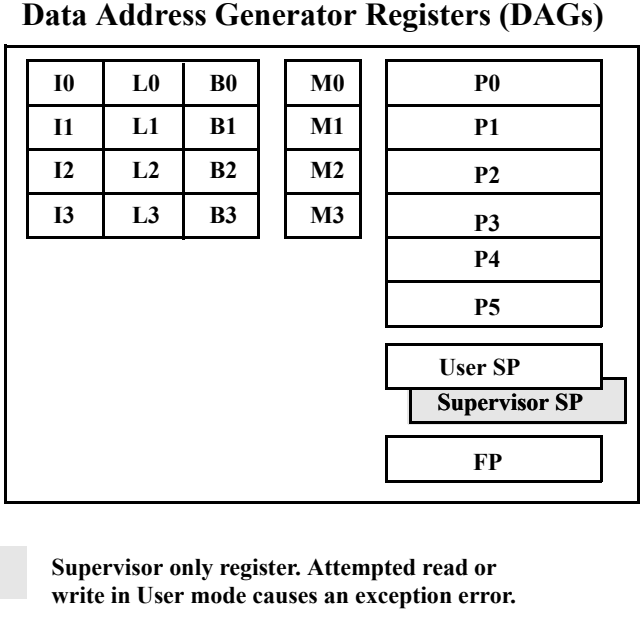

Figure 5-1. DAG Registers

Addressing With DAGs

The DAGs can generate an address that is incremented by a value or by a register. In post-modify addressing, the DAG outputs the I-register value unchanged; then the DAG adds an M-register or immediate value to the I-register.

In indexed addressing, the DAG adds a small offset to the value in the P-register, but does not update the P-register with this new value, thus providing an offset for that particular memory access.

The Blackfin processor is byte addressed. All data accesses must be aligned to the data size. In other words, a 32-bit fetch must be aligned to 32 bits, but an 8-bit store can be aligned to any byte. Depending on the type of data used, increments and decrements to the DAG registers can be by 1, 2, or 4 to match the 8-, 16-, or 32-bit accesses.

For example, consider the following instruction:

```
R0 = [ P3++ ];
```

This instruction fetches a 32-bit word, pointed to by the value in `P3`, and places it in `R0`. It then post-increments `P3` by *four*, maintaining alignment with the 32-bit access.

```
R0.L = W [ I3++ ];
```

This instruction fetches a 16-bit word, pointed to by the value in `I3`, and places it in the low half of the destination register, `R0.L`. It then post-increments `I3` by *two*, maintaining alignment with the 16-bit access.

```
R0 = B [ P3++ ] (Z) ;
```

This instruction fetches an 8-bit word, pointed to by the value in `P3`, and places it in the destination register, `R0`. It then post-increments `P3` by *one*, maintaining alignment with the 8-bit access. The byte value may be zero extended (as shown) or sign extended into the 32-bit data register.

Instructions using Index registers use an M-register or a small immediate value (+/– 2 or 4) as the modifier. Instructions using Pointer registers use

a small immediate value or another P-register as the modifier. For details, see Table 5-3 on page 5-17.

Frame and Stack Pointers

In many respects, the Frame and Stack Pointer registers perform like the other P-registers, `P[5:0]`. They can act as general pointers in any of the load/store instructions, for example, `R1 = B[SP] (Z)`. However, `FP` and `SP` have additional functionality.

The Stack Pointer registers include:

- a User Stack Pointer (`USP` in Supervisor mode, `SP` in User mode)
- a Supervisor Stack Pointer (`SP` in Supervisor mode)

The User Stack Pointer register and the Supervisor Stack Pointer register are accessed using the register alias `SP`. Depending on the current processor operating mode, only one of these registers is active and accessible as `SP`:

- In User mode, any reference to `SP` (for example, stack pop `R0 = [SP++];`) implicitly uses the `USP` as the effective address.
- In Supervisor mode, the same reference to `SP` (for example, `R0 = [SP++];`) implicitly uses the Supervisor Stack Pointer as the effective address.
- To manipulate the User Stack Pointer for code running in Supervisor mode, use the register alias `USP`. When in Supervisor mode, a register move from `USP` (for example, `R0 = USP ;`) moves the current User Stack Pointer into `R0`. The register alias `USP` can only be used in Supervisor mode.

Some load/store instructions use FP and SP exclusively:

- FP-indexed load/store, which extends the addressing range for 16-bit encoded load/stores
- Stack push/pop instructions, including those for pushing and popping multiple registers
- Link/unlink instructions, which control stack frame space and manage the Frame Pointer register (FP) for that space

Addressing Circular Buffers

The DAGs support addressing circular buffers. Circular buffers are a range of addresses containing data that the DAG steps through repeatedly, wrapping around to repeat stepping through the same range of addresses in a circular pattern.

The DAGs use four types of DAG registers for addressing circular buffers. For circular buffering, the registers operate this way:

- The Index (I) register contains the value that the DAG outputs on the address bus.
- The Modify (M) register contains the post-modify amount (positive or negative) that the DAG adds to the I-register at the end of each memory access.
- Any M-register can be used with any I-register. The modify value can also be an immediate value instead of an M-register. The size of the modify value must be less than or equal to the length (L-register) of the circular buffer.
- The Length (L) register sets the size of the circular buffer and the address range through which the DAG circulates the I-register.

- L is positive and cannot have a value greater than $2^{32} - 1$. If an L-register's value is zero, its circular buffer operation is disabled.
- The Base (B) register or the B-register plus the L-register is the value with which the DAG compares the modified I-register value after each access.

To address a circular buffer, the DAG steps the Index pointer (I-register) through the buffer values, post-modifying and updating the index on each access with a positive or negative modify value from the M-register.

If the Index pointer falls outside the buffer range, the DAG subtracts the length of the buffer (L-register) from the value or adds the length of the buffer to the value, wrapping the Index pointer back to a point inside the buffer.

The starting address that the DAG wraps around is called the buffer's base address (B-register). There are no restrictions on the value of the base address for circular buffers that contains 8-bit data. Circular buffers that contain 16- and 32-bit data must be 16-bit aligned and 32-bit aligned, respectively. Exceptions can be made for video operations. For more information, see "Memory Address Alignment" on page 5-13. Circular buffering uses post-modify addressing.

As seen in Figure 5-2, on the first post-modify access to the buffer, the DAG outputs the I-register value on the address bus, then modifies the address by adding the modify value.

- If the updated index value is within the buffer length, the DAG writes the value to the I-register.
- If the updated index value exceeds the buffer length, the DAG subtracts (for a positive modify value) or adds (for a negative modify value) the L-register value before writing the updated index value to the I-register.

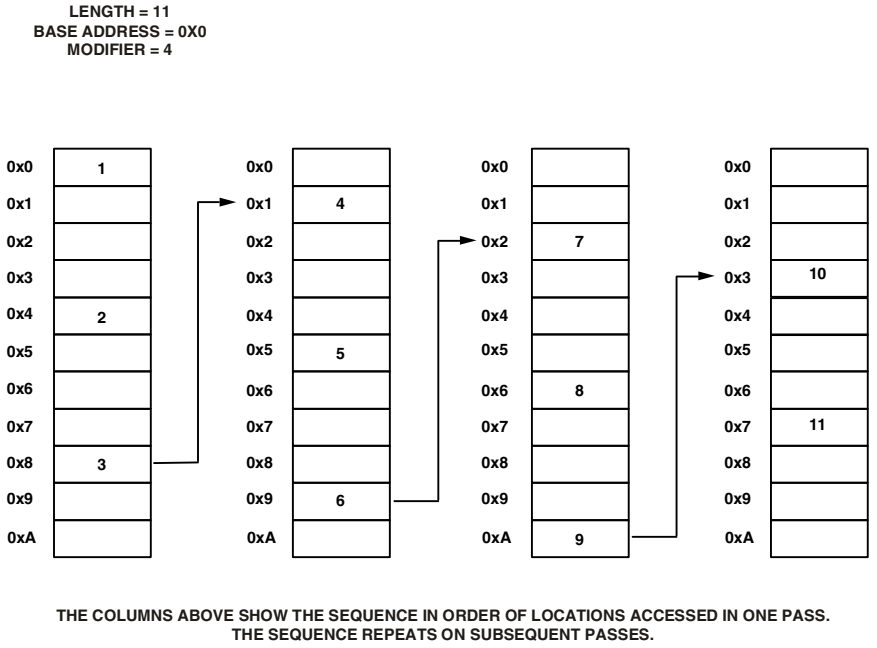

Figure 5-2. Circular Data Buffers

In equation form, these post-modify and wraparound operations work as follows, shown for "I+M" operations.

- If M is positive:

  $I_{new} = I_{old} + M$
  if $I_{old} + M <$ buffer base + length (end of buffer)

  $I_{new} = I_{old} + M - L$
  if $I_{old} + M \geq$ buffer base + length (end of buffer)

- If M is negative:

  $I_{new} = I_{old} + M$
  if $I_{old} + M \geq$ buffer base (start of buffer)

  $I_{new} = I_{old} + M + L$
  if $I_{old} + M <$ buffer base (start of buffer)

Addressing With Bit-reversed Addresses

To obtain results in sequential order, programs need bit-reversed carry addressing for some algorithms, particularly Fast Fourier Transform (FFT) calculations. To satisfy the requirements of these algorithms, the DAG's bit-reversed addressing feature permits repeatedly subdividing data sequences and storing this data in bit-reversed order. For detailed information about bit-reversed addressing, see the Modify-Increment instruction in the *Blackfin Processor Programming Reference*.

Indexed Addressing With Index and Pointer Registers

Indexed addressing uses the value in the Index or Pointer register as an effective address.

This instruction can load or store 16- or 32-bit values. The default is a 32-bit transfer. If a 16-bit transfer is required, then the *W* designator is used to preface the load or store.

For example:

```
R0 = [ I2 ];
```

loads a 32-bit value from an address pointed to by `I2` and stores it in the destination register `R0`.

```
R0.H = W [ I2 ];
```

loads a 16-bit value from an address pointed to by `I2` and stores it in the 16-bit destination register `R0.H`.

```
[ P1 ] = R0 ;
```

is an example of a 32-bit store operation.

Pointer registers can be used for 8-bit loads and stores.

For example:

```
B [ P1++ ] = R0 ;
```

stores the 8-bit value from the `R0` register in the address pointed to by the `P1` register, then increments the `P1` register.

Auto-Increment and Auto-Decrement Addressing

Auto-increment addressing updates the Pointer and Index registers after the access. The amount of increment depends on the word size. An access of 32-bit words results in an update of the Pointer by 4. A 16-bit word access updates the Pointer by 2, and an access of an 8-bit word updates the Pointer by 1. Both 8- and 16-bit read operations may specify either to sign-extend or zero-extend the contents into the destination register. Pointer registers may be used for 8-, 16-, and 32-bit accesses while Index registers may be used only for 16- and 32-bit accesses.

For example:

```
R0 = W [ P1++ ](Z);
```

loads a 16-bit word into a 32-bit destination register from an address pointed to by the `P1` Pointer register. The Pointer is then incremented by 2 and the word is zero extended to fill the 32-bit destination register.

Auto-decrement works the same way by decrementing the address after the access.

For example:

```
R0 = [ I2-- ] ;
```

loads a 32-bit value into the destination register and decrements the Index register by 4.

Pre-Modify Stack Pointer Addressing

The only pre-modify instruction in the processor uses the Stack Pointer register, `SP`. The address in `SP` is decremented by 4 and then used as an effective address for the store. The instruction `[--SP ] = R0` ; is used for stack push operations and can support only a 32-bit word transfer.

Indexed Addressing With Immediate Offset

Indexed addressing allows programs to obtain values from data tables, with reference to the base of that table. The Pointer register is modified by the immediate field and then used as the effective address. The value of the Pointer register is not updated.

Alignment exceptions are triggered when a final address is unaligned.

For example, if `P1 = 0x13`, then `[P1 + 0x11]` would effectively be equal to `[0x24]`, which is aligned for all accesses.

Post-Modify Addressing

Post-modify addressing uses the value in the Index or Pointer registers as the effective address and then modifies it by the contents of another register. Pointer registers are modified by other Pointer registers. Index registers are modified by a Modify register. This instruction does not support the Pointer registers as a destination register, nor does it support byte-addressing.

For example:

```
R5 = [ P1++P2 ] ;
```

loads a 32-bit value into the R5 register, found in the memory location pointed to by the P1 register.

The value in the P2 register is then added to the value in the P1 register.

For example:

```
R2 = W [ P4++P5 ] (Z) ;
```

loads a 16-bit word into the low half of the destination register R2 and zero-extends it to 32 bits. The value of the pointer P4 is incremented by the value of the pointer P5.

For example:

```
R2 = [ I2++M1 ];
```

loads a 32-bit word into the destination register R2. The value in the Index register, I2, is updated by the value in the Modify register, M1.

Modifying DAG and Pointer Registers

The DAGs support operations that modify an address value in an Index register without outputting an address. The operation, address-modify, is useful for maintaining pointers.

The address-modify operation modifies addresses in any DAG Index and Pointer register (I[3:0], P[5:0], FP, SP) without accessing memory. If the Index register's corresponding B- and L-registers are set up for circular buffering, the address-modify operation performs the specified buffer wraparound (if needed).

The syntax is similar to post-modify addressing (index += modifier). For Index registers, an M-register is used as the modifier. For Pointer registers, another P-register is used as the modifier.

Consider the example, `I1 += M2 ;`

This instruction adds `M2` to `I1` and updates `I1` with the new value.

Memory Address Alignment

The Blackfin processors require proper memory alignment to be maintained for the data size being accessed. Unless exceptions are disabled, violations of memory alignment cause an alignment exception. Some instructions—for example, many of the Video ALU instructions—automatically disable alignment exceptions because the data may not be properly aligned when stored in memory. Alignment exceptions may be disabled by issuing the `DISALGNEXPT` instruction in parallel with a load/store operation.

Normally, the memory system requires two address alignments:

- 32-bit word load/stores are accessed on four-byte boundaries, meaning the two least significant bits of the address are b#00.
- 16-bit word load/stores are accessed on two-byte boundaries, meaning the least significant bit of the address must be b#0.

Table 5-1 summarizes the types of transfers and transfer sizes supported by the addressing modes.

Table 5-1. Types of Transfers Supported and Transfer Sizes

Addressing Mode	Types of Transfers Supported	Transfer Sizes
Auto-increment Auto-decrement Indirect Indexed	To and from Data Registers	LOADS: 32-bit word 16-bit, zero extended half word 16-bit, sign extended half word 8-bit, zero extended byte 8-bit, sign extended byte STORES: 32-bit word 16-bit half word 8-bit byte
	To and from Pointer Registers	LOAD: 32-bit word STORE: 32-bit word
Post-increment	To and from Data Registers	LOADS: 32-bit word 16-bit half word to Data Register high half 16-bit half word to Data Register low half 16-bit, zero extended half word 16-bit, sign extended half word STORES: 32-bit word 16-bit half word from Data Register high half 16-bit half word from Data Register low half

Be careful when using the `DISALGNEXPT` instruction, because it disables automatic detection of memory alignment errors. The `DISALGNEXPT` instruction only affects misaligned loads that use I-register indirect addressing. Misaligned loads using P-register addressing will still cause an exception.

Table 5-2 summarizes the addressing modes. In the table, an asterisk (*) indicates the processor supports the addressing mode.

Table 5-2. Addressing Modes

	32-bit word	16-bit half-word	8-bit byte	Sign/zero extend	Data Register	Pointer register	Data Register Half
P Auto-inc [P0++]	*	*	*	*	*	*	
P Auto-dec [P0--]	*	*	*	*	*	*	
P Indirect [P0]	*	*	*	*	*	*	*
P Indexed [P0+im]	*	*	*	*	*	*	
FP indexed [FP+im]	*				*	*	
P Post-inc [P0++P1]	*	*		*	*		*
I Auto-inc [I0++]	*	*			*		*
I Auto-dec [I0--]	*	*			*		*
I Indirect [I0]	*	*			*		*
I Post-inc [I0++M0]	*				*		

DAG Instruction Summary

Table 5-3 lists the DAG instructions. For more information on assembly language syntax, see the *Blackfin Processor Programming Reference*. In Table 5-3, note the meaning of these symbols:

- Dreg denotes any Data Register File register.
- Dreg_lo denotes the lower 16 bits of any Data Register File register.
- Dreg_hi denotes the upper 16 bits of any Data Register File register.
- Preg denotes any Pointer register, `FP`, or `SP` register.
- Ireg denotes any DAG Index register.
- Mreg denotes any DAG Modify register.
- W denotes a 16-bit wide value.
- B denotes an 8-bit wide value.
- immA denotes a signed, A-bits wide, immediate value.
- uimmAmB denotes an unsigned, A-bits wide, immediate value that is an even multiple of B.
- Z denotes the zero-extension qualifier.
- X denotes the sign-extension qualifier.
- BREV denotes the bit-reversal qualifier.

The *Blackfin Processor Programming Reference* more fully describes the options that may be applied to these instructions and the sizes of immediate fields.

DAG instructions do not affect the ASTAT Status flags.

Table 5-3. DAG Instruction Summary

Instruction
Preg = [Preg] ;
Preg = [Preg ++] ;
Preg = [Preg --] ;
Preg = [Preg + uimm6m4] ;
Preg = [Preg + uimm17m4] ;
Preg = [Preg – uimm17m4] ;
Preg = [FP – uimm7m4] ;
Dreg = [Preg] ;
Dreg = [Preg ++] ;
Dreg = [Preg --] ;
Dreg = [Preg + uimm6m4] ;
Dreg = [Preg + uimm17m4] ;
Dreg = [Preg – uimm17m4] ;
Dreg = [Preg ++ Preg] ;
Dreg = [FP – uimm7m4] ;
Dreg = [Ireg] ;
Dreg = [Ireg ++] ;
Dreg = [Ireg --] ;
Dreg = [Ireg ++ Mreg] ;
Dreg =W [Preg] (Z) ;
Dreg =W [Preg ++] (Z) ;
Dreg =W [Preg --] (Z) ;
Dreg =W [Preg + uimm5m2] (Z) ;
Dreg =W [Preg + uimm16m2] (Z) ;

Table 5-3. DAG Instruction Summary (Cont'd)

Instruction
Dreg =W [Preg – uimm16m2] (Z) ;
Dreg =W [Preg ++ Preg] (Z) ;
Dreg = W [Preg] (X) ;
Dreg = W [Preg ++] (X) ;
Dreg = W [Preg --] (X) ;
Dreg =W [Preg + uimm5m2] (X) ;
Dreg =W [Preg + uimm16m2] (X) ;
Dreg =W [Preg – uimm16m2] (X) ;
Dreg =W [Preg ++ Preg] (X) ;
Dreg_hi = W [Ireg] ;
Dreg_hi = W [Ireg ++] ;
Dreg_hi = W [Ireg --] ;
Dreg_hi = W [Preg] ;
Dreg_hi = W [Preg ++ Preg] ;
Dreg_lo = W [Ireg] ;
Dreg_lo = W [Ireg ++] ;
Dreg_lo = W [Ireg --] ;
Dreg_lo = W [Preg] ;
Dreg_lo = W [Preg ++ Preg] ;
Dreg = B [Preg] (Z) ;
Dreg = B [Preg ++] (Z) ;
Dreg = B [Preg --] (Z) ;
Dreg = B [Preg + uimm15] (Z) ;
Dreg = B [Preg – uimm15] (Z) ;
Dreg = B [Preg] (X) ;
Dreg = B [Preg ++] (X) ;

Table 5-3. DAG Instruction Summary (Cont'd)

Instruction
Dreg = B [Preg --] (X) ;
Dreg = B [Preg + uimm15] (X) ;
Dreg = B [Preg – uimm15] (X) ;
[Preg] = Preg ;
[Preg ++] = Preg ;
[Preg --] = Preg ;
[Preg + uimm6m4] = Preg ;
[Preg + uimm17m4] = Preg ;
[Preg – uimm17m4] = Preg ;
[FP – uimm7m4] = Preg ;
[Preg] = Dreg ;
[Preg ++] = Dreg ;
[Preg --] = Dreg ;
[Preg + uimm6m4] = Dreg ;
[Preg + uimm17m4] = Dreg ;
[Preg – uimm17m4] = Dreg ;
[Preg ++ Preg] = Dreg ;
[FP – uimm7m4] = Dreg ;
[Ireg] = Dreg ;
[Ireg ++] = Dreg ;
[Ireg --] = Dreg ;
[Ireg ++ Mreg] = Dreg ;
W [Ireg] = Dreg_hi ;
W [Ireg ++] = Dreg_hi ;
W [Ireg --] = Dreg_hi ;
W [Preg] = Dreg_hi ;

Table 5-3. DAG Instruction Summary (Cont'd)

Instruction
W [Preg ++ Preg] = Dreg_hi ;
W [Ireg] = Dreg_lo ;
W [Ireg ++] = Dreg_lo ;
W [Ireg --] = Dreg_lo ;
W [Preg] = Dreg_lo ;
W [Preg] = Dreg ;
W [Preg ++] = Dreg ;
W [Preg --] = Dreg ;
W [Preg + uimm5m2] = Dreg ;
W [Preg + uimm16m2] = Dreg ;
W [Preg – uimm16m2] = Dreg ;
W [Preg ++ Preg] = Dreg_lo ;
B [Preg] = Dreg ;
B [Preg ++] = Dreg ;
B [Preg --] = Dreg ;
B [Preg + uimm15] = Dreg ;
B [Preg – uimm15] = Dreg ;
Preg = imm7 (X) ;
Preg = imm16 (X) ;
Preg += Preg (BREV) ;
Ireg += Mreg (BREV) ;
Preg = Preg << 2 ;
Preg = Preg >> 2 ;
Preg = Preg >> 1 ;
Preg = Preg + Preg << 1 ;
Preg = Preg + Preg << 2 ;

Table 5-3. DAG Instruction Summary (Cont'd)

Instruction
Preg –= Preg ;
Ireg –= Mreg ;

6 MEMORY

The ADSP-BF561 has dual cores and each core supports a hierarchical memory model with different performance and size parameters, depending on the memory location within the hierarchy. Level 1 (L1) memories are located on the chip and faster, whereas Level 2 (L2) memories are memory systems farther from the core, typically with longer access latencies. The faster L1 memories, which include instruction, data and scratchpad memory as part of the Blackfin core are accessed in a single-cycle. The L2 memories, which include an on-chip SRAM and off-chip synchronous and asynchronous devices, provide much higher memory space with higher latency.

Memory Architecture

The ADSP-BF561 has a unified 4GB address range that spans a combination of on-chip and off-chip memory and memory-mapped I/O resources. Of this range, 272 MB of address space is dedicated to internal, on-chip resources. The ADSP-BF561 populates portions of this internal memory space with:

- A set of L1 and L2 Static Random Access Memories (SRAM)
- A set of memory-mapped registers (MMRs)
- A boot Read-Only Memory (ROM)

A portion of the internal L1 SRAM can also be configured to run as cache. The ADSP-BF561 also provides support for an external memory space that includes asynchronous memory space and synchronous DRAM (SDRAM) space. See Chapter 16, “External Bus Interface Unit”, for a detailed discussion of each of these memory regions and the controllers that support them.

The diagram in Figure 6-1 provides an overview of the ADSP-BF561 system memory map. Note that the architecture does not define a separate I/O space. All resources are mapped through the flat 32-bit address space. The memory is byte-addressable.

As shown in Figure 6-1, total on-chip memory for the each DSP core occupies 100KB, as follows:

- 32KB of instruction memory:
    - 16KB of instruction SRAM
    - 16KB of instruction cache/SRAM, lockable by Way or line
- 64KB of data memory:
    - 32KB of data cache/SRAM
    - 32KB of SRAM
- 4KB of data scratch pad SRAM

An on-chip SRAM provides 128KB of L2 memory space. For systems using some or all ADSP-BF561 L1 memory as cache, the on-chip L2 SRAM memory can help provide deterministic, bounded-memory access times.

The upper portion of internal memory space is allocated to the core and system MMRs of the ADSP-BF561. Accesses to this area are allowed only when the processor is in Supervisor mode or Emulation mode (see Chapter 3, “Operating Modes and States”.

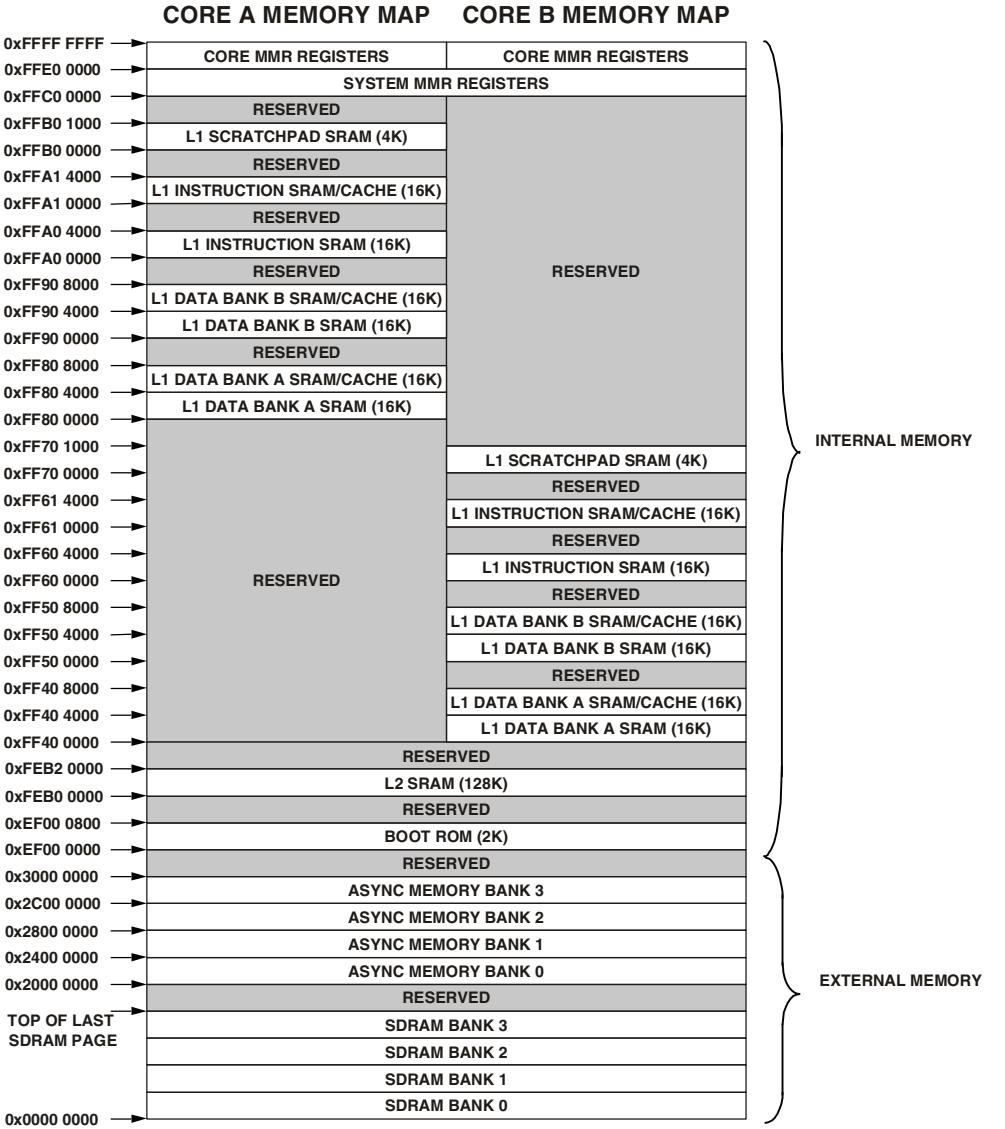

Figure 6-1. Memory Map

The lowest 2KB of internal memory space is occupied by the boot ROM of the ADSP-BF561. Depending on the booting option selected, the appropriate boot program is executed from this memory space when the ADSP-BF561 is reset (see "Booting Methods" on page 3-19.)

Within the external memory map, four banks of asynchronous memory space and four banks of SDRAM memory are available. Each of the asynchronous banks is 64MB and each of the synchronous banks can be configured 16-128 MB.

Internal Memory

The ADSP-BF561 L1 memory system performance provides high bandwidth and low latency. Because SRAMs provide deterministic access time and very high throughput, DSP systems have traditionally achieved performance improvements by providing fast SRAM on chip. The ADSP-BF561 supports this memory architecture for applications that require direct control over access time.

The addition of instruction and data caches (SRAMs with cache control hardware) provides both high performance and a simple programming model. Caches eliminate the need to explicitly manage data movement into and out of the L1 memories. Code can be ported to or developed for the ADSP-BF561 quickly without requiring performance optimization for the memory organization.

Each core's L1 memory provides:

- A modified Harvard architecture, allowing up to four core memory accesses per clock cycle (one 64-bit instruction fetch, and two 32-bit data loads and one pipelined 32-bit data store.)
- Simultaneous system DMA, cache maintenance, and core accesses
- SRAM access at processor clock rate (CCLK) for critical DSP algorithms and fast context switching

- Instruction and data cache options for microcontroller code, excellent High-Level Language (HLL) support, and ease of programming cache control instructions, such as `PREFETCH` and `FLUSH`
- Memory protection

The L1 memories operate at the core clock frequency (`CCLK`).

Overview of L1 Instruction SRAM

The 32KB L1 instruction SRAM consists of a dedicated 16KB SRAM plus an additional 16KB bank which can be configured as either SRAM or cache. The upper 16KB, L1 Instruction Memory can be configured as a 4-Way set-associative cache (see Figure 6-4 on page 6-13). Consequently, instructions can be brought into four different ways of cache, decreasing the frequency of cache-line replacements and increasing overall performance. When the upper 16KB of L1 memory is configured as a cache, individual ways or lines of L1 instruction cache can be locked down, allowing further control over the location of time-critical code. The cache-locking concept is explained further in "Instruction Cache Locking by Way" on page 6-18. When configured as SRAM, each of the two 16KB banks of memory is broken into 4KB sub-banks which can be independently accessed by the processor and DMA. For more information about L1 instruction SRAM, see "L1 Instruction SRAM" on page 6-9.

Overview of L1 Data SRAM

Each core on the ADSP-BF561 provides two 32KB, L1 data SRAM banks (Data Bank A and Data Bank B). Each data bank has a dedicated lower 16KB SRAM bank plus an additional upper 16KB bank which can be configured as SRAM or cache.

When configured as cache, the upper 16KB bank in each L1 data bank is a 2-Way set-associative structure. This provides two separate locations that

can hold cached data, decreasing the rate of cache-line replacements and increasing overall performance.

If configured as SRAM, each of the two upper 16KB banks of memory is broken into four 4KB sub-banks which can be independently accessed by the processor and DMA. For more information about L1 data SRAM, see "L1 Data SRAM" on page 6-26.

Overview of Scratchpad Data SRAM

The processor provides a dedicated 4K byte bank of scratchpad data SRAM. The scratchpad is independent of the configuration of the other L1 memory banks and cannot be configured as cache or targeted by DMA. Typical applications use the scratchpad data memory where speed is critical. For example, the User and Supervisor stacks should be mapped to the scratchpad memory for the fastest context switching during interrupt handling.

The L1 memories operate at the core clock frequency (CCLK).

Scratchpad data SRAM cannot be accessed by the DMA controller.

Overview of On-Chip L2 Memory

The on-chip Level 2 (L2) memory provides 128KB of low latency, high-bandwidth capacity. This memory system is referred to as on-chip L2 because it forms an on-chip memory hierarchy with L1 memory. On-chip L2 memory provides much more capacity than L1 memory, but the latency is higher. The on-chip L2 memory is SRAM and can not be configured as cache. It is capable of storing both instructions and data. The L1 caches can be configured to cache instructions and data located in the on-chip L2 memory.

L1 Instruction Memory

L1 Instruction Memory consists of a combination of dedicated SRAM and banks which can be configured as SRAM or cache. For the 16KB bank that can be either cache or SRAM, control bits in the IMEM_CONTROL register can be used to organize all four sub-banks of the L1 Instruction Memory as any of the following:

- A simple SRAM
- A 4-Way, set-associative instruction cache
- A cache with as many as four locked Ways

L1 Instruction Memory can be used only to store instructions.

Instruction Memory Control Register (IMEM_CONTROL)

The Instruction Memory Control register (IMEM_CONTROL) contains control bits for the L1 Instruction Memory. By default after reset, cache and Cacheability Protection Lookaside Buffer (CPLB) address checking is disabled (see "L1 Instruction Cache" on page 6-11).

When the LRUPRIORST bit is set to 1, the cached states of all CPLB_LRUPRIO bits (see "ICPLB Data Registers (ICPLB_DATAx)" on page 6-56) are cleared. This simultaneously forces all cached lines to be of equal (low) importance. Cache replacement policy is based first on line importance indicated by the cached states of the CPLB_LRUPRIO bits, and then on LRU (least recently used). See "Instruction Cache Locking by Line" on page 6-17 for complete details. This bit must be 0 to allow the state of the CPLB_LRUPRIO bits to be stored when new lines are cached.

The ILOC[3:0] bits provide a useful feature only after code has been manually loaded into cache. See "Instruction Cache Locking by Way" on page 6-18. These bits specify which Ways to remove from the cache replacement policy. This has the effect of locking code present in

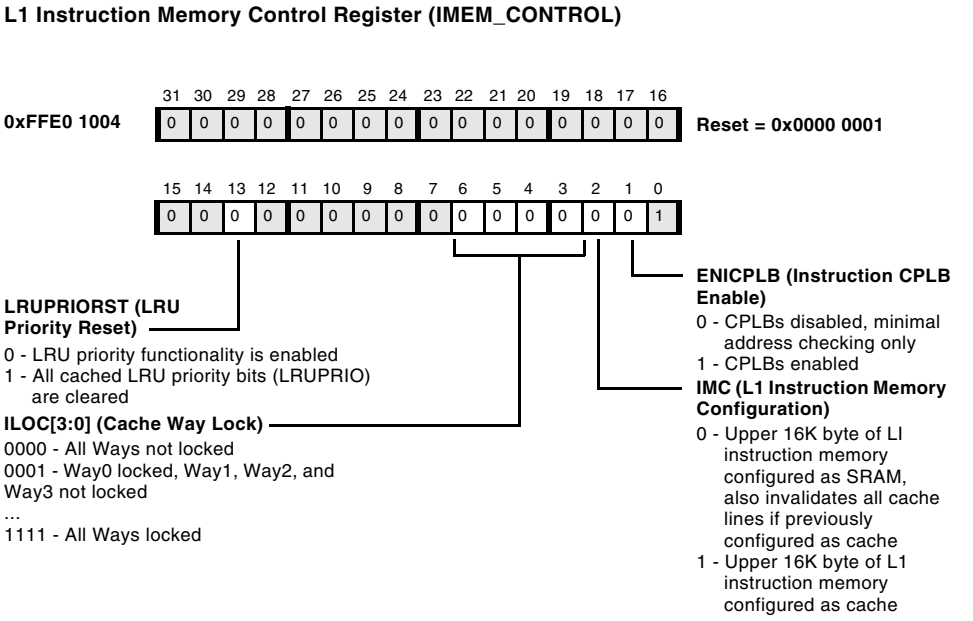

Figure 6-2. L1 Instruction Memory Control Register

nonparticipating Ways. Code in nonparticipating Ways can still be removed from the cache using an `IFLUSH` instruction. If an `ILOC[3:0]` bit is 0, the corresponding Way is not locked and that Way participates in cache replacement policy. If an `ILOC[3:0]` bit is 1, the corresponding Way is locked and does not participate in cache replacement policy.

The `IMC` bit reserves a portion of L1 instruction SRAM to serve as cache. Note reserving memory to serve as cache will not alone enable L2 memory accesses to be cached. CPLBs must also be enabled using the `EN_ICPLB` bit and the CPLB descriptors (`ICPLB_DATAx` and `ICPLB_ADDRx` registers) must specify desired memory pages as cache-enabled.

Instruction CPLBs are disabled by default after reset. When disabled, only minimal address checking is performed by the L1 memory interface. This

minimal checking generates an exception to the processor whenever it attempts to fetch an instruction from:

- Reserved (non populated) L1 instruction memory space
- L1 data memory space
- MMR space

CPLBs must be disabled using this bit prior to updating their descriptors (DCPLB_DATAx and DCPLB_ADDRx registers). Note since load store ordering is weak (see "Ordering of Loads and Stores" on page 6-68), disabling of CPLBs should be proceeded by a CSYNC.

When enabling or disabling cache or CPLBs, immediately follow the write to IMEM_CONTROL with a SSYNC to ensure proper behavior.

To ensure proper behavior and future compatibility, all reserved bits in this register must be set to 0 whenever this register is written.

L1 Instruction SRAM

The ADSP-BF561 core reads the instruction memory through the 64-bit-wide instruction-fetch bus. All addresses from this bus are 64-bit aligned. Each instruction fetch can return any combination of 16-, 32- or 64-bit instructions (for example, four 16-bit instructions, two 16-bit instructions and one 32-bit instruction, or one 64-bit instruction).

The DAGs, which are described in Chapter 5, cannot access L1 Instruction Memory directly. A DAG reference to instruction memory SRAM space generates an exception (see "Exceptions" on page 4-43).

Write access to the L1 instruction SRAM memory must be made through the 64-bit-wide system DMA port. Because the SRAM is implemented as a collection of single-ported sub-banks, the instruction memory is effectively dual-ported. Provided that system and core accesses do not access

the same 32-bit polarity (address bits 2 match) of the same sub-bank, effective dual-porting of the instruction memory is achieved. If both system and core attempt to access the same 32-bit polarity (address bits 2 match) of the same bank, the system DMA controller has priority over the core instruction fetch.

Table 6-1. L1 Instruction Memory Sub-Banks

Memory Subbank	Memory Start Location, Core A	Memory Start Location, Core B
0	0xFFA0 0000	0xFF60 0000
1	0xFFA0 1000	0xFF60 1000
2	0xFFA0 2000	0xFF60 2000
3	0xFFA0 3000	0xFF60 3000
4	0xFFA1 0000	0xFF61 0000
5	0xFFA1 1000	0xFF61 1000
6	0xFFA1 2000	0xFF61 2000
7	0xFFA1 3000	0xFF61 3000

Before changing the configuration state, be sure to flush the cache or move all modified data from the SRAM, if so configured.

Figure 6-3 describes the bank architecture of the L1 Instruction Memory. As the figure shows, each 16KB bank is made up of four 4KB sub-banks.

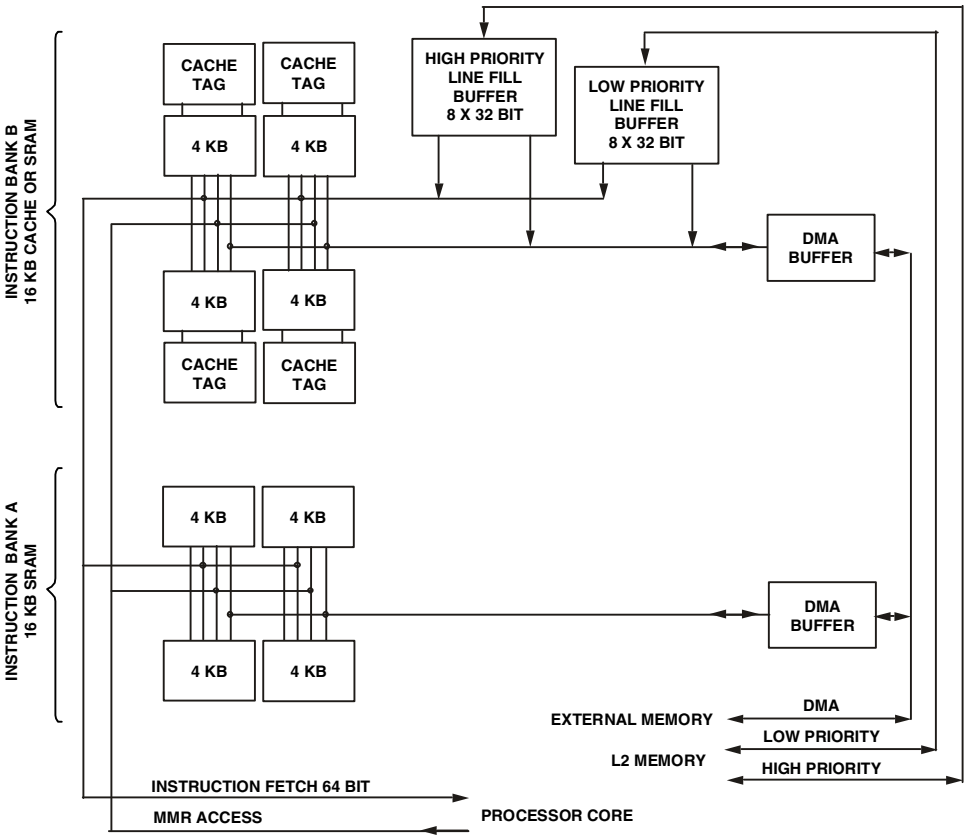

Figure 6-3. L1 Instruction Memory Bank Architecture

L1 Instruction Cache

The L1 Instruction Memory may also be configured as a flexible, 4-Way set-associative instruction 16KB cache. To improve the average access latency for critical code sections, each Way of the cache can be locked

independently. When the memory is configured as cache, it cannot be accessed directly.

When cache is enabled, only memory pages specified as cacheable by Cacheability Protection Lookaside Buffers (CPLBs) will be cached. When CPLBs are enabled, any memory location that is accessed must have an associated page definition available, or a CPLB exception will be generated. CPLBs are described in "Memory Protection and Properties" on page 6-47.

Figure 6-4 shows the overall Blackfin instruction cache organization.

Cache Lines

As shown in Figure 6-4, the cache consists of a collection of cache lines. Each cache line is made up of a *tag* component and a *data* component:

- The tag component incorporates a 20-bit address tag, least-recently-used (LRU) bits, a Valid bit, and a Line Lock bit.
- The data component is made up of four 64-bit words of instruction data.

  The tag and data components of cache lines are stored in the tag and data memory arrays, respectively.

The address tag consists of the upper 18 bits plus bits 11 and 10 of the physical address. Bits 12 and 13 of the physical address are not part of the address tag. Instead, these bits are used to identify the 4KB memory sub-bank targeted for the access.

The LRU bits are part of an LRU algorithm used to determine which cache line should be replaced if a cache miss occurs.

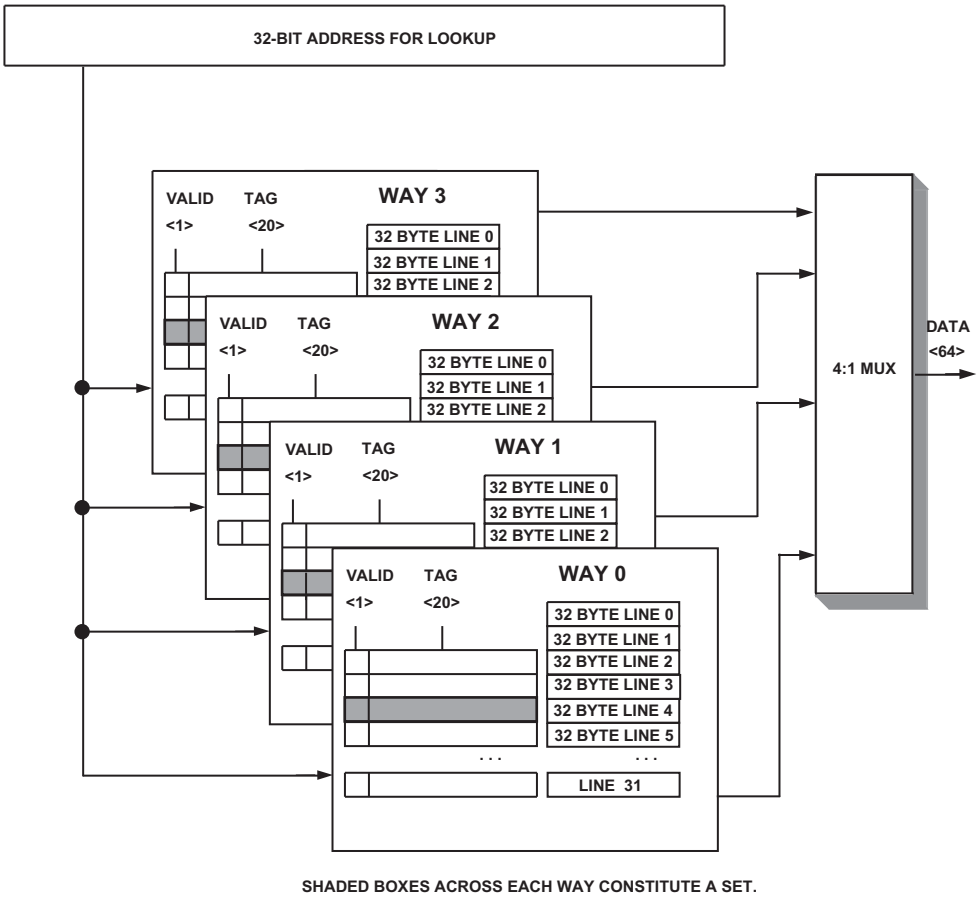

Figure 6-4. Blackfin Instruction Cache Organization

The Valid bit indicates the state of a cache line. A cache line is always valid or invalid:

- Invalid cache lines have their Valid bit cleared, indicating the line will be ignored during an address-tag compare operation.
- Valid cache lines have their Valid bit set, indicating the line contains valid instruction/data that is consistent with the source memory.

The tag and data components of a cache line are illustrated in Figure 6-5.

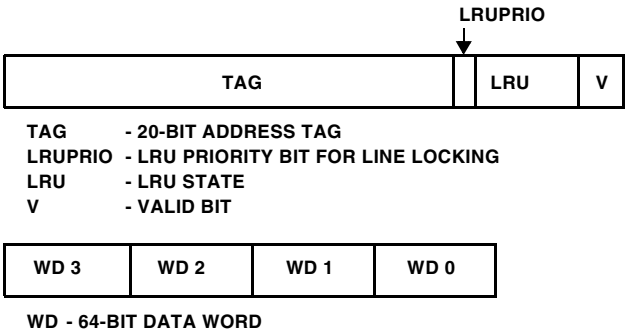

Figure 6-5. Cache Line – Tag and Data Portions

Cache Hits and Misses

A cache hit occurs when the address for an instruction-fetch request from the core matches a valid entry in the cache. Specifically, a cache hit is determined by comparing the upper 18 bits and bits 11 and 10 of the instruction-fetch address to the address tags of valid lines currently stored in a cache set. The cache set is selected, using bits 9 through 5 of the instruction-fetch address. If the address-tag compare operation results in a match, a cache hit occurs. If the address-tag compare operation does not result in a match, a cache miss occurs.

When a cache hit occurs, the target 64-bit instruction word is first sent to the Instruction Alignment Unit where it is stored in one of two 64-bit instruction buffers.

When a cache miss occurs, the instruction memory unit generates a cache line-fill access to retrieve the missing cache line from memory that is external to the core. The address for the on-chip L2 or external memory access is the address of the target instruction word. When a cache miss occurs, the core halts until the target instruction word is returned from on-chip L2 or external memory.

Cache-Line Fills

A cache-line fill consists of fetching 32 bytes of data from memory. The operation starts when the instruction memory unit requests a line-read data transfer (that is. a burst of four 64-bit words of data) on its on-chip L2 or external read-data port. The address for the read transfer is the address of the target instruction word. When responding to a line-read request from the instruction memory unit, the on-chip L2 or external memory returns the target instruction word first. After it has returned the target instruction word, the next three words are fetched in sequential address order. This fetch will wrap around if necessary, as shown in Table 6-2.

Table 6-2. Cache-Line Word-Fetching Order

Target Word	Fetching Order for Next Three Words
WD0	WD0, WD1, WD2, WD3
WD1	WD1, WD2, WD3, WD0
WD2	WD2, WD3, WD0, WD1
WD3	WD3, WD0, WD1, WD2

Line-Fill Buffer

As the new cache line is retrieved from on-chip L2 or external memory, each 64-bit word is buffered into one of two four-entry line-fill buffer before it is written to a 4KB memory bank within L1 memory. The line-fill buffer allows the core to access the data from the new cache line as the line is being retrieved from on-chip L2 or external memory, rather than having to wait until the line has been written into the cache.

Two separate line fill buffers are provided to allow a load from slow external memory to continue without causing jumps to higher speed on-chip L2 memory to stall. Which line buffer is used, is determined by the CPLB_MEMLEV bit in the memory pages CPLB's. See "Memory Protection and Properties" on page 6-47.

Cache-Line Replacement

When the instruction memory unit is configured as cache, bits 9 through 5 of the instruction fetch address are used as the index to select the cache set for the tag-address compare operation. If the tag-address compare operation results in a cache miss, the Valid bits for the selected set are examined by a cache-line replacement unit to determine the entry to use for the new cache line, that is, whether to use Way0, Way1, Way2, or Way3 (see Figure 6-4 on page 6-13).

The cache-line replacement unit first checks for invalid entries (that is, entries having its Valid bit cleared). If only a single invalid entry is found, that entry is selected for the new cache line. If multiple invalid entries are found, the replacement entry for the new cache line is selected based on the following priority:

- Way0 first
- Way1 next
- Way2 next
- Way3 last

For example:

- If Way3 is invalid and Ways0, 1, 2 are valid, Way3 is selected for the new cache line.
- If Ways0 and 1 are invalid and Ways2 and 3 are valid, Way0 is selected for the new cache line.
- If Ways2 and 3 are invalid and Ways0 and 1 are valid, Way2 is selected for the new cache line.

When no invalid entries are found, the cache replacement logic uses an LRU algorithm.

Instruction Cache Management

The system DMA controller and the core DAGs cannot access the instruction cache directly. By a combination of instructions and the use of core MMRs, it is possible to initialize the instruction tag and data arrays indirectly and provide a mechanism for instruction cache test, initialization, and debug.

The coherency of instruction cache must be explicitly managed. To accomplish this and ensure that the instruction cache fetches the latest version of any modified instruction space, invalidate instruction cache line entries, as required.

For more information, see “Instruction Cache Invalidation” on page 6-19.

Instruction Cache Locking by Line

The `CPLB_LRUPRIO` bits in the `ICPLB_DATAx` registers (see “Memory Protection and Properties” on page 6-47) are used to enhance control over which code remains resident in the instruction cache. When a cache line is filled, the state of this bit is stored along with the line’s tag. It is then used in conjunction with the LRU (least recently used) policy to determine which Way is victimized when all cache Ways are occupied when a new

cacheable line is fetched. This bit indicates that a line is of either “low” or “high” importance. In a modified LRU policy, a high can replace a low, but a low cannot replace a high. If all Ways are occupied by highs, an otherwise cacheable low will still be fetched for the core, but will not be cached. Fetched highs seek to replace unoccupied Ways first, then least recently used lows next, and finally other highs using the LRU policy. Lows can only replace unoccupied Ways or other lows, and do so using the LRU policy. If *all* previously cached highs ever become less important, they may be simultaneously transformed into lows by writing to the `LRUPRIRST` bit in the `IMEM_CONTROL` register (on page 6-7).

Instruction Cache Locking by Way

The instruction cache has four independent lock bits (`ILOC[3:0]`) that control each of the four Ways of the instruction cache. When the cache is enabled, L1 Instruction Memory has four Ways available. Setting the lock bit for a specific Way prevents that way from participating in the LRU replacement policy. Thus, a cached instruction, with its way locked, can only be removed using an IFLUSH instruction, or “backdoor” MMR assisted manipulation of the tag array.

An example sequence is provided below to demonstrate how to lock down Way0:

- If the code of interest may already reside in the instruction cache, invalidate the entire cache first (for an example, see “Instruction Cache Invalidation” on page 6-19).
- Disable interrupts, if required, to prevent Interrupt Service Routines (ISRs) from potentially corrupting the locked cache.
- Set the locks for the other Ways of the cache by setting `ILOC[3:1]`. Only Way0 of the instruction cache can now be replaced by new code.

- Execute the code of interest. Any cacheable exceptions, such as exit code, traversed by this code execution are also locked into the instruction cache.
- Upon exit of the critical code, clear `ILOC[3:1]`, and set `ILOC[0]`. The critical code (and the instructions which set `ILOC[0]`), are now locked into Way0.
- Re-enable interrupts, if required.

If all four Ways of the cache are locked, then further allocation into the cache is prevented.

Instruction Cache Invalidation

The instruction cache can be invalidated by an address, cache line, or a complete cache. The `IFLUSH` instruction can explicitly invalidate cache lines based on their line addresses. The target address of the instruction is generated from the P-registers. Because the instruction cache should not contain modified (dirty) data, the cache line is simply invalidated.

In the following example, the `P2` register contains the address of a valid memory location. If this address has been brought into cache, the corresponding cache line is invalidated after the execution of this instruction.

Example of `ICACHE` instruction:

```
iflush [ p2 ] ;   /* Invalidate cache line containing address
that P2 points to */
```

Because the `IFLUSH` instruction is used to invalidate a specific address in the ADSP-BF561 memory map, it is impractical to use this instruction to invalidate an entire bank of cache. A second, faster technique can be used to invalidate an entire cache bank directly. This second technique directly invalidates Valid bits by setting the Invalid bit of each cache line to the invalid state. To implement this technique, additional MMRs (`ITEST_COMMAND` and `ITEST_DATA[1:0]`) are available to allow arbitrary

read/write of all cache entries directly. This method is explained in the next section.

For invalidating the complete instruction cache, a third method is available. By clearing the IMC bit in the IMEM_CONTROL register (see Figure 6-2 on page 6-8), all valid bits in the instruction cache are set to the invalid state. A second write to the IMEM_CONTROL register to set the IMC bit then configures the instruction memory as cache again. An SSYNC should be run before invalidating the cache and a CSYNC should be inserted after each of these operations.

Instruction Test Registers

The Instruction Test registers allow arbitrary read/write of all L1 cache entries directly. They make it possible to initialize the instruction tag and data arrays and to provide a mechanism for instruction cache test, initialization, and debug.

When the Instruction Test Command register (ITEST_COMMAND) is used, the L1 cache data or tag arrays are accessed, and data is transferred through the Instruction Test Data registers (ITEST_DATA[1:0]). The ITEST_DATAx registers contain either the 64-bit data that the access is to write to or the 64-bit data that was read during the access. The lower 32 bits are stored in the ITEST_DATA[0] register, and the upper 32 bits are stored in the ITEST_DATA[1] register. When the tag arrays are accessed, ITEST_DATA[0] is used. Graphical representations of the ITEST registers begin with Figure 6-6.

The following figures describe the ITEST registers:

- Figure 6-6 on page 6-21
- Figure 6-7 on page 6-22
- Figure 6-8 on page 6-23

Access to these registers is possible only in Supervisor or Emulation mode. When writing to ITEST registers, always write to the ITEST_DATAx registers first, then the ITEST_COMMAND register. When reading from ITEST registers, reverse the sequence—read the ITEST_COMMAND register first, then the ITEST_DATAx registers.

ITEST_COMMAND Register

When the Instruction Test Command register (ITEST_COMMAND) is written to, the L1 cache data or tag arrays are accessed, and the data is transferred through the Instruction Test Data registers (ITEST_DATA[1:0]).

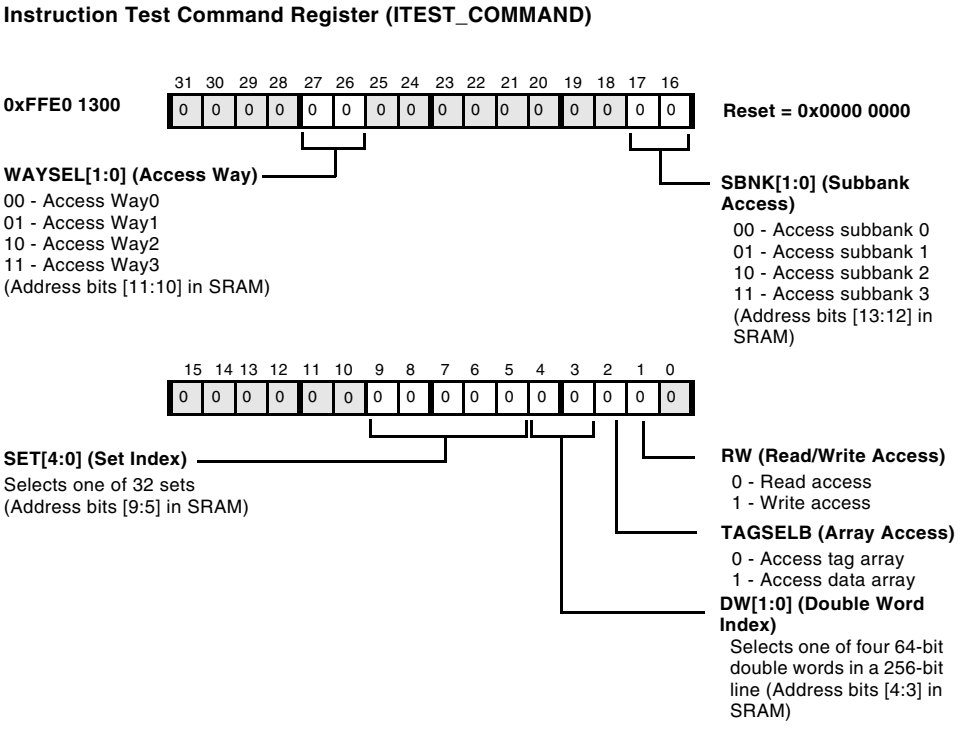

Figure 6-6. Instruction Test Command Register

ITEST_DATA1 Register

Instruction Test Data registers (ITEST_DATA[1:0]) are used to access L1 cache data arrays (see Figure 6-7). They contain either the 64-bit data that the access is to write to or the 64-bit data that the access is to read from. The Instruction Test Data 1 register (ITEST_DATA1) stores the upper 32 bits.

Instruction Test Data 1 Register (ITEST_DATA1)

Used to access L1 cache data arrays and tag arrays. When accessing a data array, stores the upper 32 bits of 64-bit words of instruction data to be written to or read from by the access. See "Cache Lines" on page 6-12.

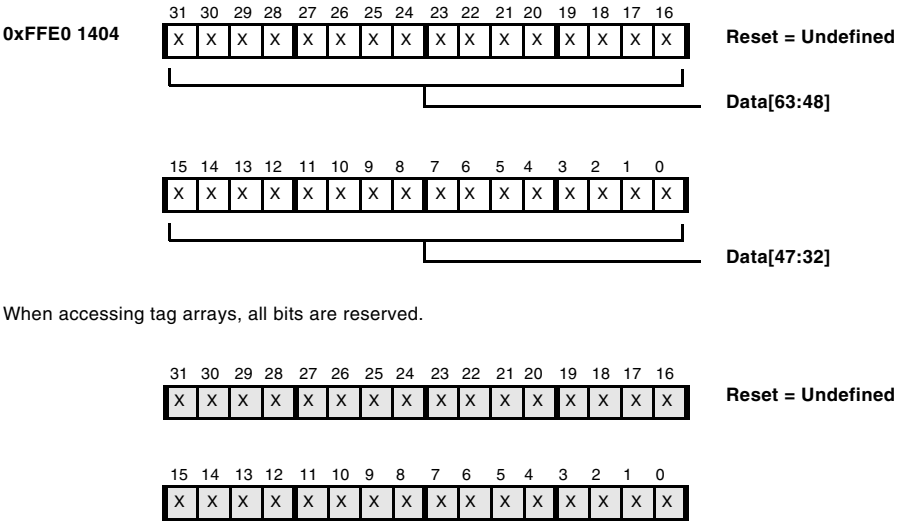

Figure 6-7. Instruction Test Data 1 Register

ITEST_DATA0 Register

The Instruction Test Data 0 register (ITEST_DATA0) stores the lower 32 bits of the 64-bit data to be written to or read from by the access (see Figure 6-8).

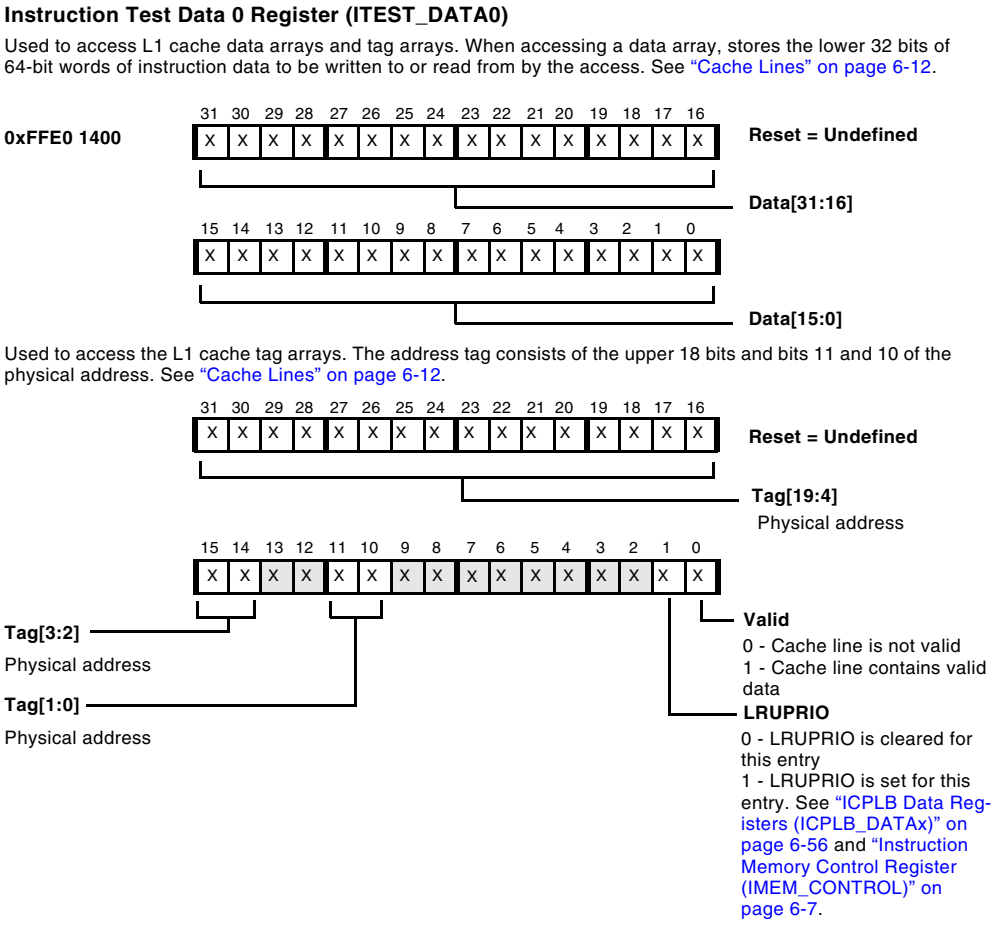

Figure 6-8. Instruction Test Data 0 Register

The ITEST_DATA0 register is also used to access tag arrays. This register also contains the Valid and Dirty bits, which indicate the state of the cache line.

L1 Data Memory

The L1 data SRAM/cache is constructed from single-ported subsections, but organized to reduce the likelihood of access collisions. This organization results in apparent multi-ported behavior. When there are no collisions, this L1 data traffic could occur in a single core clock cycle:

- Two 32-bit DAG loads
- One pipelined 32-bit DAG store
- One 64-bit DMA IO
- One 64-bit cache fill/victim access

L1 Data Memory can be used only to store data.

Data Memory Control Register (DMEM_CONTROL)

The Data Memory Control register (DMEM_CONTROL) contains control bits for the L1 Data Memory.

The PORT_PREF1 bit selects the data port used to process DAG1 non-cacheable L2 fetches. Cacheable fetches are always processed by the data port physically associated with the targeted cache memory. Steering DAG0, DAG1, and cache traffic to different ports optimizes performance by keeping the queue to L2 memory full.

The PORT_PREF0 bit selects the data port used to process DAG0 non-cacheable L2 fetches. Cacheable fetches are always processed by the data port physically associated with the targeted cache memory. Steering

DAG0, DAG1, and cache traffic to different ports optimizes performance by keeping the queue to L2 memory full.

For optimal performance with dual DAG reads, DAG0 and DAG1 should be configured for different ports. For example, if PORT_PREF0 is configured as 1, then PORT_PREF1 should be programmed to 0.

The DCBS bit provides some control over which addresses alias into the same set. This bit can be used to affect which addresses tend to remain resident in cache by avoiding victimization of repetitively used sets. It has no affect unless both Data Bank A and Data Bank B are serving as cache (bits DMC[1:0] in this register are set to 11).

The ENDCPLB bit is used to enable/disable the 16 Cacheability Protection Lookaside Buffers (CPLBs) used for data (see "L1 Data Cache" on page 6-30). Data CPLBs are disabled by default after reset. When disabled, only minimal address checking is performed by the L1 memory interface. This minimal checking generates an exception when the processor:

- Addresses nonexistent (reserved) L1 memory space
- Attempts to perform a nonaligned memory access
- Attempts to access MMR space either using DAG1 or when in User mode
- Attempts to write the on-chip boot ROM

CPLBs must be disabled using this bit prior to updating their descriptors (registers DCPLB_DATAx and DCPLB_ADDRx). Note that since load store ordering is weak (see "Ordering of Loads and Stores" on page 6-68), disabling CPLBs should be preceded by a CSYNC instruction, and enabling CPLBs should be followed by a CSYNC instruction in order to ensure predictable behavior.

When enabling or disabling cache or CPLBs, immediately follow the write to DMEM_CONTROL with a SSYNC to ensure proper behavior.

By default after reset, all L1 Data Memory serves as SRAM. The DMC[1:0] bits can be used to reserve portions of this memory to serve as cache instead. Reserving memory to serve as cache does not enable L2 memory accesses to be cached. To do this, CPLBs must also be enabled (using the ENDCPLB bit) and CPLB descriptors (registers DCPLB_DATAx and DCPLB_ADDRx) must specify chosen memory pages as cache-enabled.

By default after reset, cache and CPLB address checking is disabled.

To ensure proper behavior and future compatibility, all reserved bits in this register must be set to 0 whenever this register is written.

L1 Data SRAM

Accesses to SRAM do not collide unless they are to the same 32-bit word polarity (address bits 2 match), the same 4KB subbank (address bits 13 and 12 match), the same 16KB half bank (address bits 16 match), and the same bank (address bits 21 and 20 match). When an address collision is detected, access is nominally granted first to the DAGs, then to the store buffer, and finally to the DMA and cache fill/victim traffic. To ensure adequate DMA bandwidth, DMA is given highest priority if it has been blocked for more than 16 sequential core clock cycles, or if a second DMA I/O is queued before the first DMA I/O is processed.

Data Memory Control Register (DMEM_CONTROL)

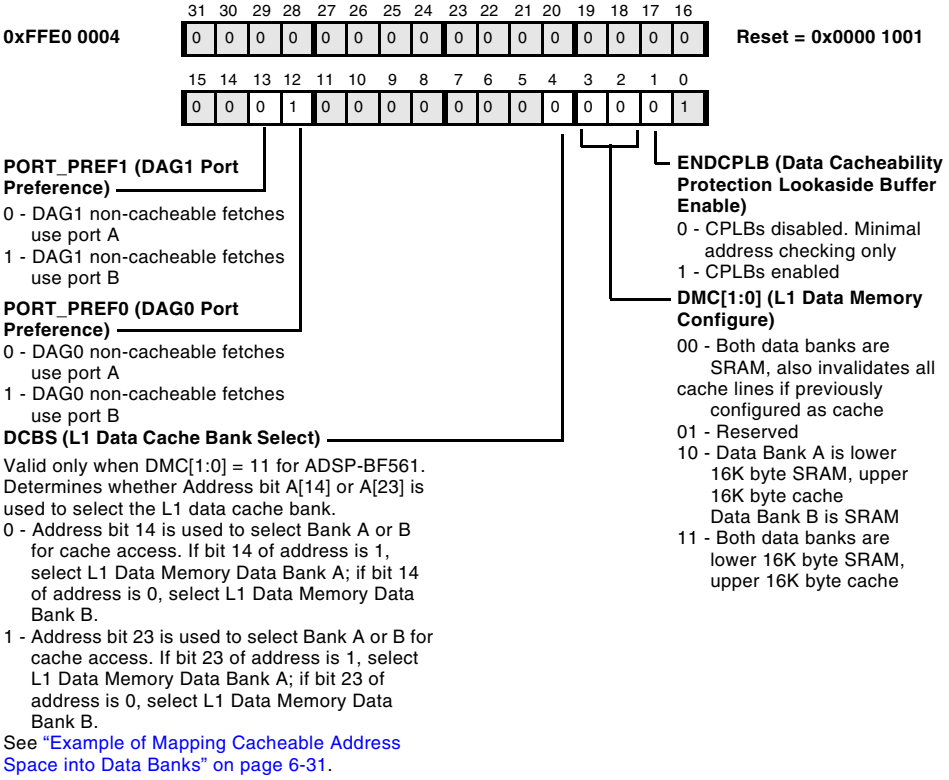

Figure 6-9. L1 Data Memory Control Register

Table 6-3 shows how the subbank organization is mapped into memory.

Table 6-3. L1 Data Memory SRAM Subbank Start Addresses

Memory Bank and Subbank	Core A	Core B
Data Bank A, Subbank 0	0xFF80 0000	0xFF40 0000
Data Bank A, Subbank 1	0xFF80 1000	0xFF40 1000
Data Bank A, Subbank 2	0xFF80 2000	0xFF40 2000
Data Bank A, Subbank 3	0xFF80 3000	0xFF40 3000
Data Bank A, Subbank 4	0xFF80 4000	0xFF40 4000
Data Bank A, Subbank 5	0xFF80 5000	0xFF40 5000
Data Bank A, Subbank 6	0xFF80 6000	0xFF40 6000
Data Bank A, Subbank 7	0xFF80 7000	0xFF40 7000
Data Bank B, Subbank 0	0xFF90 0000	0xFF50 0000
Data Bank B, Subbank 1	0xFF90 1000	0xFF50 1000
Data Bank B, Subbank 2	0xFF90 2000	0xFF50 2000
Data Bank B, Subbank 3	0xFF90 3000	0xFF50 3000
Data Bank B, Subbank 4	0xFF90 4000	0xFF50 4000
Data Bank B, Subbank 5	0xFF90 5000	0xFF50 5000
Data Bank B, Subbank 6	0xFF90 6000	0xFF50 6000
Data Bank B, Subbank 7	0xFF90 7000	0xFF50 7000

Figure 6-10 shows the L1 Data Memory architecture.

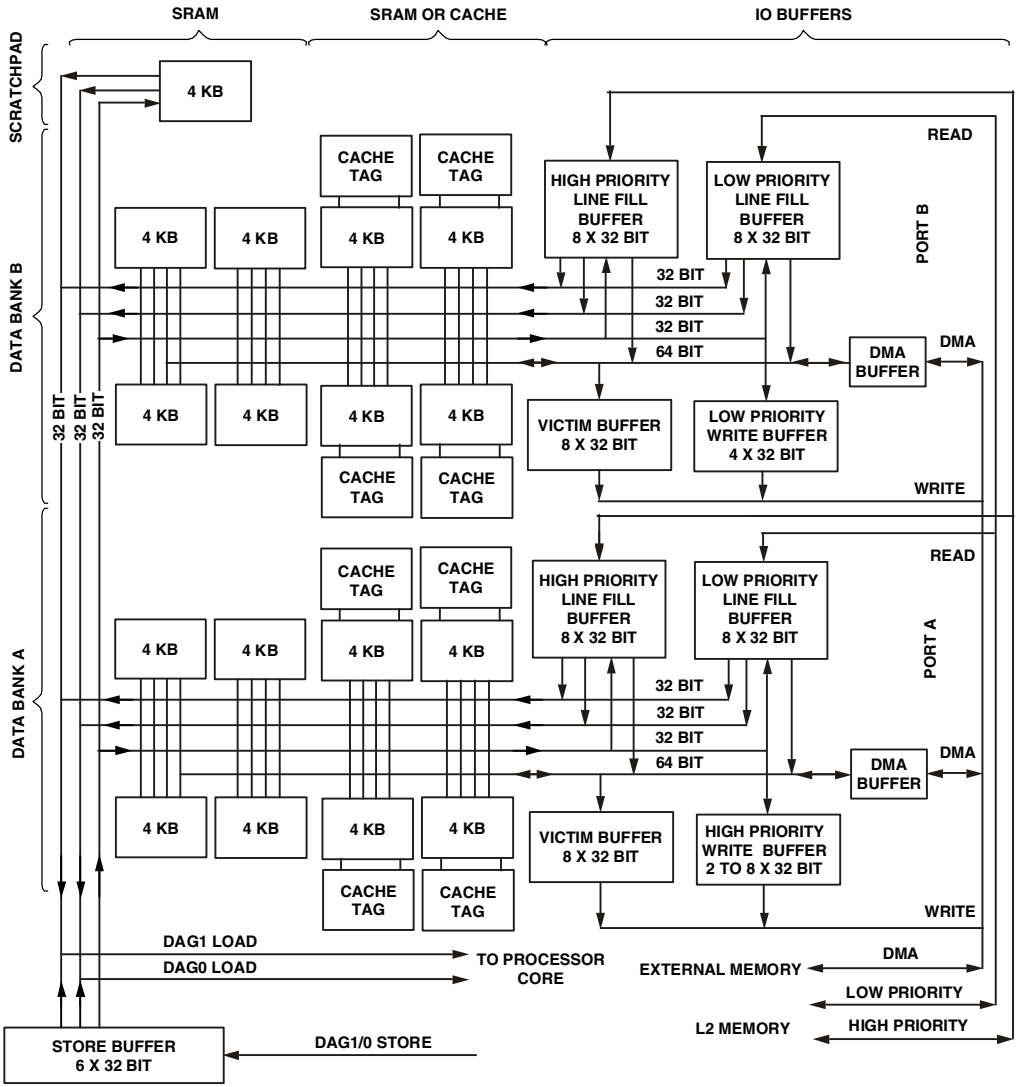

Figure 6-10. L1 Data Memory Architecture for ADSP-BF561

L1 Data Cache

For definitions of cache terminology, see "Terminology" on page 6-75.

When data cache is enabled (controlled by bits DMC[1:0] in the DMEM_CONTROL register), either 16KB of Data Bank A or 16KB of both Data Bank A and Data Bank B can be set to serve as cache. For the ADSP-BF561, the upper 16KB is used. Unlike instruction cache, which is 4-Way set associative, data cache is 2-Way set associative. When two banks are available and enabled as cache, additional sets rather than Ways are created. When both Data Bank A and Data Bank B have memory serving as cache, the DCBS bit in the DMEM_CONTROL register may be used to control which half of all address space is handled by which bank of cache memory. The DCBS bit selects either address bit 14 or 23 to steer traffic between the cache banks. This provides some control over which addresses alias into the same set. It may therefore be used to affect which addresses tend to remain resident in cache by avoiding victimization of repetitively used sets.

Accesses to cache do not collide unless they are to the same 4KB subbank (address bits 13 and 12 match), the same half bank (address bits 16 match), and to the same bank (address bits 21 and 20 match). Cache has less apparent multi-ported behavior due to the overhead in maintaining tags. When cache addresses collide, access is granted first to the DTEST register accesses, then to the store buffer, and finally to cache fill/victim traffic.

Three different cache modes are available.

- Write-through with cache line allocation only on reads
- Write-through with cache line allocation on both reads and writes
- Write-back which allocates cache lines on both reads and writes

Cache mode is selected by the DCPLB descriptors (see "Memory Protection and Properties" on page 6-47). Any combination of these cache modes can

be used simultaneously since cache mode is selectable for each memory page independently.

If cache is enabled (controlled by bits DMC[1:0] in the DMEM_CONTROL register), data CPLBs should also be enabled (controlled by ENDCPLB bit in the DMEM_CONTROL register). Only memory pages specified as cacheable by data CPLBs will be cached. The default behavior when data CPLBs are disabled is for nothing to be cached.

Erroneous behavior can result when MMR space is configured as cacheable by data CPLBs, or when data banks serving as L1 SRAM are configured as cacheable by data CPLBs.

Example of Mapping Cacheable Address Space into Data Banks

An example of how the cacheable address space maps into two data banks follows.

When both banks are configured as cache on the ADSP-BF561, they operate as two independent, 16KB, 2-Way set associative caches that can be independently mapped into the Blackfin processor address space.

If both data banks are configured as cache, the DCBS bit in the DMEM_CONTROL register designates Address bit A[14] or A[23] as the cache selector. Address bit A[14] or A[23] selects the cache implemented by Data Bank A or the cache implemented by Data Bank B.

- If DCBS = 0, then A[14] is part of the address index, and all addresses in which A[14] = 0 use Data Bank A. All addresses in which A[14] = 1 use Data Bank B.

  In this case, A[23] is treated as merely another bit in the address that is stored with the tag in the cache and compared for hit/miss processing by the cache.

- If DCBS = 1, then A[23] is part of the address index, and all addresses where A[23] = 0 use Data Bank A. All addresses where A[23] = 1 use Data Bank B.

  In this case, A[14] is treated as merely another bit in the address that is stored with the tag in the cache and compared for hit/miss processing by the cache.

The result of choosing DCBS = 0 or DCBS = 1 is:

- If DCBS = 0, A[14] selects Data Bank A instead of Data Bank B.

  Alternating 16KB pages of memory map into each of the two 16KB caches implemented by the two data banks. Consequently:

  - Any data in the first 16KB of memory could be stored only in Data Bank A.
  - Any data in the next address range (16KB through 32KB) – 1 could be stored only in Data Bank B.
  - Any data in the next range (32KB through 48KB) – 1 would be stored in Data Bank A.
  - Alternate mapping would continue.

  As a result, the cache operates as if it were a single, contiguous, 2-Way set associative 32KB cache. Each Way is 16KB long, and all data elements with the same first 14 bits of address index to a unique set in which up to two elements can be stored (one in each Way).

- If DCBS = 1, A[23] selects Data Bank A instead of Data Bank B.

  With DCBS = 1, the system functions more like two independent caches, each a 2-Way set associative 16KB cache. Each Bank serves an alternating set of 8 MB blocks of memory. For example, Data Bank A caches all data accesses for the first 8 MB of memory

address range. That is, every 8 MB of range vies for the two line entries (rather than every 16KB repeat). Likewise, Data Bank B caches data located above 8 MB and below 16 MB.

For example, if the application is working from a data set that is 1 MB long and located entirely in the first 8 MB of memory, it is effectively served by only half the cache, that is, by Data Bank A (a 2-Way set associative 16KB cache). In this instance, the application never derives any benefit from Data Bank B.

For most applications, it is best to operate with `DCBS = 0`. However, if the application is working from two data sets, located in two memory spaces at least 8 MB apart, closer control over how the cache maps to the data is possible. For example, if the program is doing a series of dual MAC operations in which both DAGs are accessing data on every cycle, by placing DAG0's data set in one block of memory and DAG1's data set in the other, the system can ensure that:

- DAG0 gets its data from Data Bank A for all of its accesses and
- DAG1 gets its data from Data Bank B.

This arrangement causes the core to use both data buses for cache line transfer and achieves the maximum data bandwidth between the cache and the core.

Figure 6-11 shows an example of how mapping is performed when `DCBS = 1`.

The `DCBS` selection can be changed dynamically; however, to ensure that no data is lost, first flush and invalidate the entire cache.

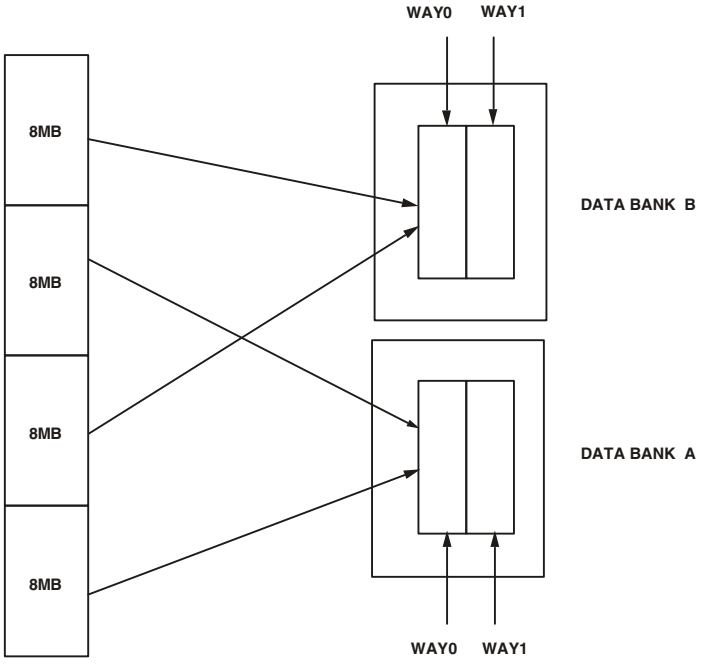

Figure 6-11. Data Cache Mapping When DCBS = 1

Data Cache Access

The Cache Controller tests the address from the DAGs against the tag bits. If the logical address is present in L1 cache, a cache hit occurs, and the data is accessed in L1. If the logical address is not present, a cache miss occurs, and the memory transaction is passed to the next level of memory via the system interface. The line index and replacement policy for the Cache Controller determines the cache tag and data space that are allocated for the data coming back from L2 memory.

A data cache line is in one of three states: invalid, exclusive (valid and clean), and modified (valid and dirty). If valid data already occupies the

allocated line and the cache is configured for write-back storage, the controller checks the state of the cache line and treats it accordingly:

- If the state of the line is exclusive (clean), the new tag and data write over the old line.
- If the state of the line is modified (dirty), then the cache contains the only valid copy of the data.
- If the line is dirty, the current contents of the cache are copied back to L2 memory before the new data is written to the cache.

The processor provides victim buffers and line fill buffers. These buffers are used if a cache load miss generates a victim cache line that should be replaced. The line fill operation goes to L2 memory. The data cache performs the line fill request to the system as critical (or requested) word first, and forwards that data to the waiting DAG as it updates the cache line. In other words, the cache performs critical word forwarding.

The data cache supports hit-under-a-store miss, and hit-under-a-prefetch miss. In other words, on a write-miss or execution of a `PREFETCH` instruction that misses the cache (and is to a cacheable region), the instruction pipeline incurs a minimum of a 4-cycle stall. Furthermore, a subsequent load or store instruction can hit in the L1 cache while the line fill completes.

Interrupts of sufficient priority (relative to the current context) cancel a stalled load instruction. Consequently, if the load operation misses the L1 Data Memory cache and generates a high latency line fill operation on the system interface, it is possible to interrupt the core, causing it to begin processing a different context. The system access to fill the cache line is not cancelled, and the data cache is updated with the new data before any further cache miss operations to the respective data bank are serviced. For more information see “Exceptions” on page 4-43.

Cache Write Method

Cache write memory operations can be implemented by using either a write-through method or a write-back method:

- For each store operation, write-through caches initiate a write to L2 memory immediately upon the write to cache.

  If the cache line is replaced or explicitly flushed by software, the contents of the cache line are invalidated rather than written back to L2 memory.

- A write-back cache does not write to L2 memory until the line is replaced by a load operation that needs the line.

  The L1 Data Memory employs a full cache line width copyback buffer on each data bank.

Write Buffers

Two separate write buffers are provided to allow stores to slow external memory to continue without causing stores to higher speed on-chip L2 memory to stall. Which buffer is used, is determined by the `CPLB_MEMLEV` bit in the data memory page's CPLBs. See "Memory Protection and Properties" on page 6-47.

These two write buffers in the L1 data memory accept all stores with each cache inhibited or store-through protection.

An `SSYNC` instruction flushes the write buffers.

Interrupt Priority Register (IPRIO) and Write Buffer Depth

The Interrupt Priority register (IPRIO) can be used to control the size of the high priority write buffer on Port A (see "L1 Data Memory Architecture for ADSP-BF561" on page 6-29).

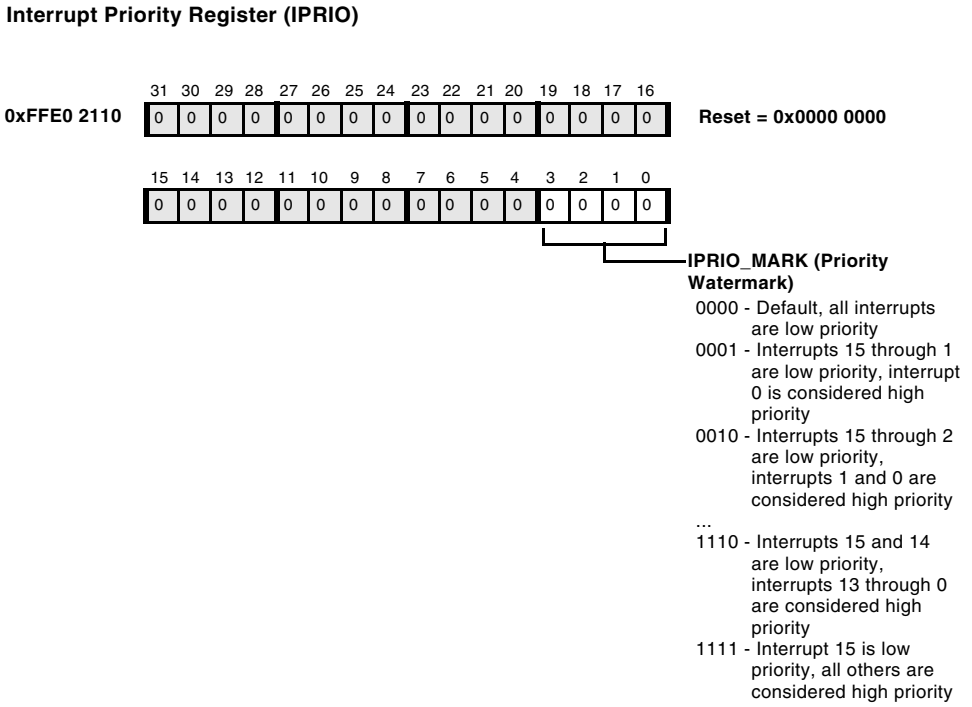

Figure 6-12. Interrupt Priority Register

The IPRIO[3:0] bits can be programmed to reflect the low priority interrupt watermark. When an interrupt occurs, causing the processor to vector from a low priority interrupt service routine to a high priority interrupt service routine, the size of the low priority write buffer increases from two to eight 32-bit words deep. This allows the interrupt service routine to run and post writes without an initial stall, in the case where the low

priority write buffer was already filled in the low priority interrupt routine. This is most useful when posted writes are to a slow external memory device. When returning from a high priority interrupt service routine to a low priority interrupt service routine or user mode, the core stalls until the write buffer has completed the necessary writes to return to a two-deep state. By default, the low priority write buffer is a fixed two-deep FIFO.

Data Cache Control Instructions

The processor defines three data cache control instructions that are accessible in User and Supervisor modes. The instructions are `PREFETCH`, `FLUSH`, and `FLUSHINV`.

- `PREFETCH` (Data Cache Prefetch) attempts to allocate a line into the L1 cache. If the prefetch hits in the cache, generates an exception, or addresses a cache inhibited region, `PREFETCH` functions like a `NOP`.
- `FLUSH` (Data Cache Flush) causes the data cache to synchronize the specified cache line with L2 memory. If the cached data line is dirty, the instruction writes the line out and marks the line clean in the data cache. If the specified data cache line is already clean or does not exist, `FLUSH` functions like a `NOP`.
- `FLUSHINV` (Data Cache Line Flush and Invalidate) causes the data cache to perform the same function as the `FLUSH` instruction and then invalidate the specified line in the cache. If the line is in the cache and dirty, the cache line is written out to L2 memory. The Valid bit in the cache line is then cleared. If the line is not in the cache, `FLUSHINV` functions like a `NOP`.

If software requires synchronization with system hardware, place an `SSYNC` instruction after the `FLUSH` instruction to ensure that the flush operation has completed. If ordering is desired to ensure that previous stores have been pushed through all the queues, place an `SSYNC` instruction before the `FLUSH`.

Data Cache Invalidation

Besides the FLUSHINV instruction, explained in the previous section, two additional methods are available to invalidate the data cache when flushing is not required. The first technique directly invalidates Valid bits by setting the Invalid bit of each cache line to the invalid state. To implement this technique, additional MMRs (DTEST_COMMAND and DTEST_DATA[1:0]) are available to allow arbitrary reads/writes of all the cache entries directly. This method is explained in the next section.

For invalidating the complete data cache, a second method is available. By clearing the DMC[1:0] bits in the DMEM_CONTROL register (see Figure 6-9 on page 6-27), all Valid bits in the data cache are set to the invalid state. A second write to the DMEM_CONTROL register to set the DMC[1:0] bits to their previous state then configures the data memory back to its previous cache/SRAM configuration. An SSYNC instruction should be run before invalidating the cache and a CSYNC instruction should be inserted after each of these operations.

Data Test Registers

Like L1 Instruction Memory, L1 Data Memory contains additional MMRs to allow arbitrary reads/writes of all cache entries directly. The registers provide a mechanism for data cache test, initialization, and debug.

When the Data Test Command register (DTEST_COMMAND) is written to, the L1 cache data or tag arrays are accessed and data is transferred through the Data Test Data registers (DTEST_DATA[1:0]). The DTEST_DATA[1:0] registers contain the 64-bit data to be written, or they contain the destination for the 64-bit data read. The lower 32 bits are stored in the DTEST_DATA[0] register and the upper 32 bits are stored in the DTEST_DATA[1] register. When the tag arrays are being accessed, then the DTEST_DATA[0] register is used.

A CSYNC instruction is required after writing the DTEST_COMMAND MMR.

These figures describe the DTEST registers.

- Figure 6-13 on page 6-41
- Figure 6-14 on page 6-42
- Figure 6-15 on page 6-43

Access to these registers is possible only in Supervisor or Emulation mode. When writing to DTEST registers, always write to the DTEST_DATA registers first, then the DTEST_COMMAND register.

Data Test Command Register (DTEST_COMMAND)

When the Data Test Command register (DTEST_COMMAND) is written to, the L1 cache data or tag arrays are accessed, and the data is transferred through the Data Test Data registers (DTEST DATA[1:0]).

The Data/Instruction Access bit allows direct access via the DTEST_COMMAND MMR to L1 instruction SRAM.

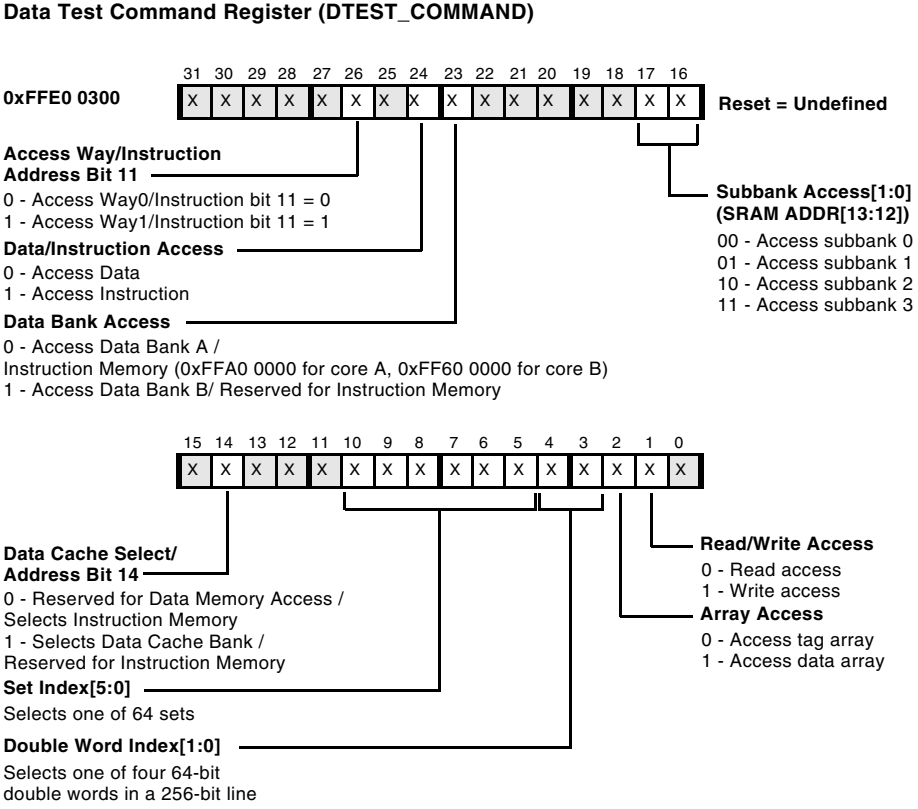

Figure 6-13. Data Test Command Register

Data Test Data 1 Register (DTEST_DATA1)

Data Test Data registers (DTEST_DATA[1:0]) contain the 64-bit data to be written, or they contain the destination for the 64-bit data read. The Data Test Data 1 register (DTEST_DATA1) stores the upper 32 bits.

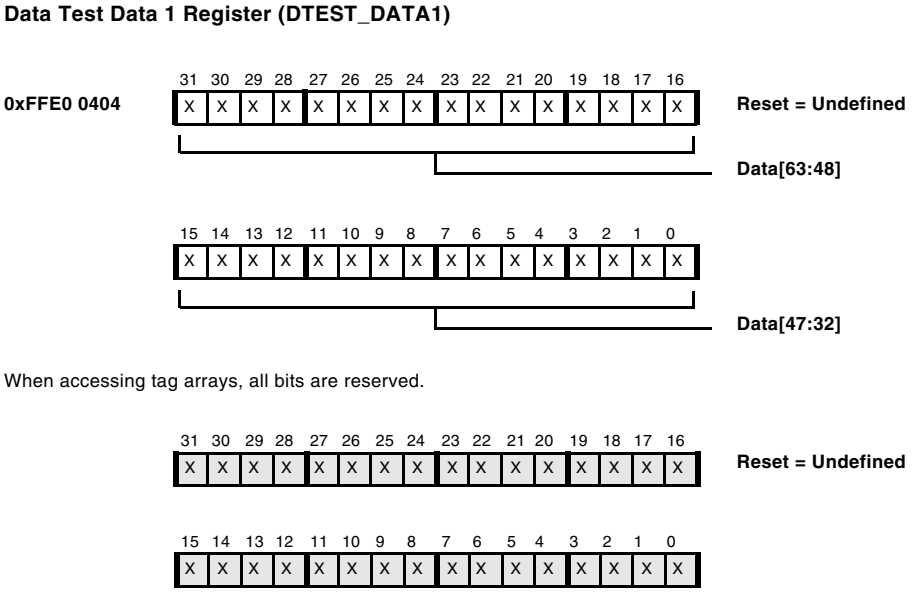

Figure 6-14. Data Test Data 1 Register

Data Test Data 0 Register (DTEST_DATA0)

The Data Test Data 0 register (DTEST_DATA0) stores the lower 32 bits of the 64-bit data to be written, or it contains the lower 32 bits of the destination for the 64-bit data read.

The DTEST_DATA0 register is also used to access the tag arrays and contains the Valid and Dirty bits, which indicate the state of the cache line.

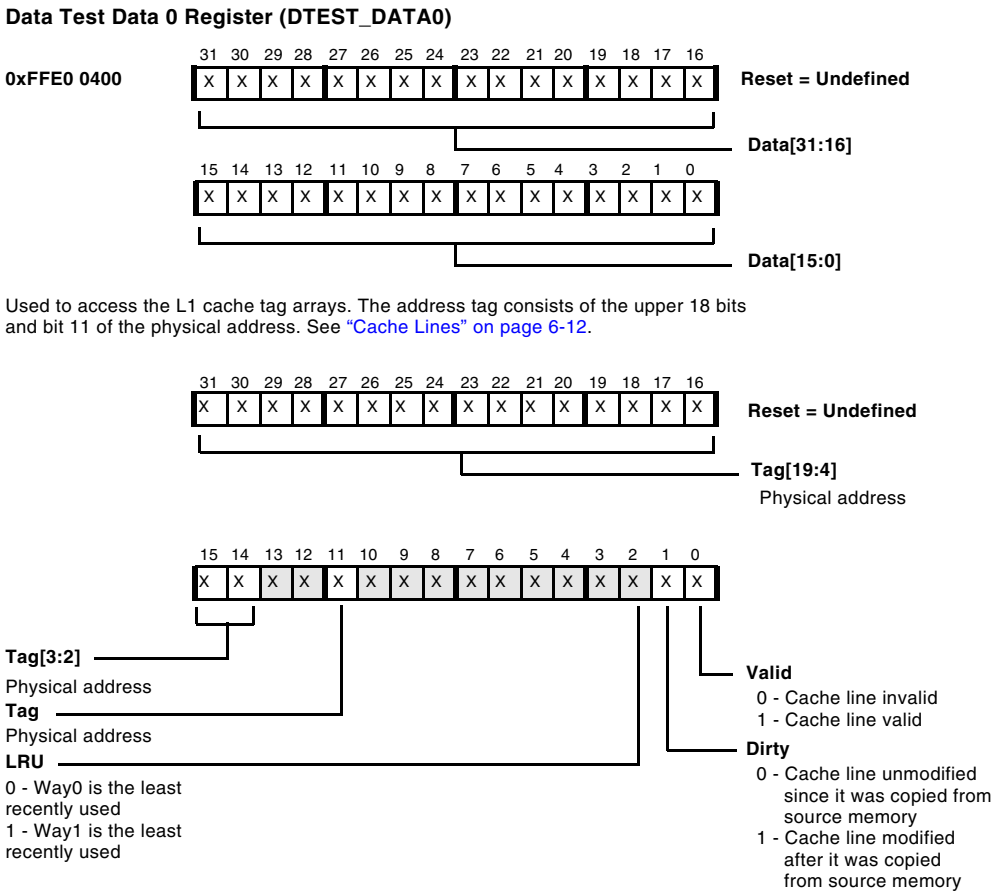

Figure 6-15. Data Test Data 0 Register

On-Chip Level 2 (L2) Memory

Configured as SRAM, the on-chip Level 2 (L2) memory of the ADSP-BF561 DSP provides 128KB of low latency, high bandwidth storage capacity. For systems that use some or all ADSP-BF561 DSP L1 memory as cache, the on-chip L2 SRAM memory system can help provide deterministic, bounded memory access times.

Simultaneous access to the multibanked, on-chip L2 memory architecture from the cores and system DMA can occur in parallel, provided that they access different banks. A fixed-priority arbitration scheme resolves conflicts. The on-chip system DMA controllers share a dedicated 32-bit data path into the L2 memory system. This interface operates at the SCLK frequency. Dedicated L2 access from the processor core is also supported.

The cores and IMDMA share a dedicated, low latency, 64-bit data path into the L2 SRAM memory. At a core clock frequency of 600 MHz, the peak data transfer rate across this interface is 4.8 GB/second.

On-Chip L2 Bank Access

The L2 memory is divided into eight separate 16K sub-banks. Two L2 access ports, a processor core port and a system port, are provided to allow concurrent access to the L2 memory, provided that the two ports access different memory sub-banks. If simultaneous access to the same memory sub-bank is attempted, collision detection logic in the L2 provides arbitration. This is a fixed priority arbiter; the DMA port always has the highest priority, unless the core is granted access to the sub-bank for a burst transfer. In this case, the L2 finishes the burst transfer before the system bus is granted access.

Latency

When cache is enabled, the bus between the core and L2 memory is fully pipelined for contiguous burst transfers. The cache line fill from on-chip

memory behaves the same for instruction and data fetches. Operations that miss the cache trigger a cache line replacement. This replacement fills one 256-bit (32-byte) line with four 64-bit reads. Under this condition, the L1 cache line fills from the L2 SRAM in `9+2+2+2=15` cycles. In other words, after nine core cycles, the first 64-bit (8-byte) fill is available for the processor. Figure 6-16 shows an example of L2 latency with cache on.

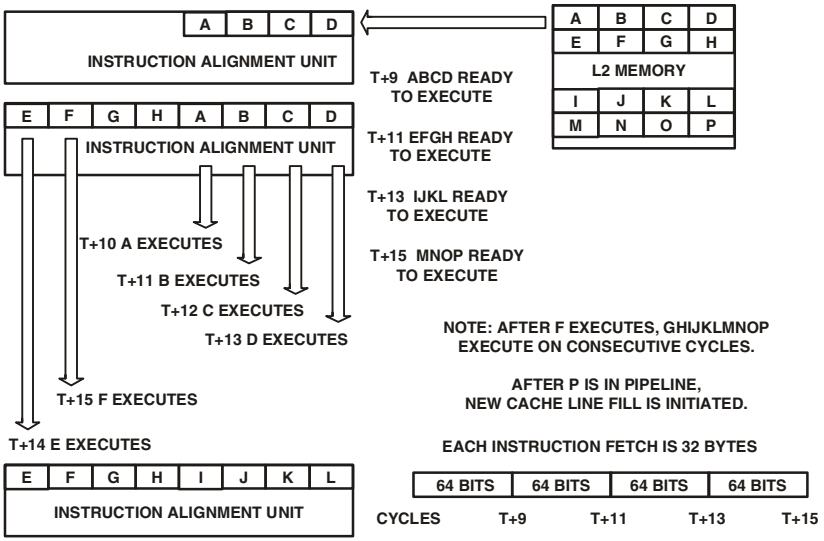

Figure 6-16. L2 Latency With Cache On

In this example, at the end of 15 core cycles, 32 bytes of instructions or data have been brought into cache and are available to the sequencer. If all the instructions contain 16 bits, sixteen instructions are brought into cache at the end of 15 cycles. In addition, the first instruction that is part of the cache-line fill executes on the tenth cycle; the second instruction executes on the eleventh cycle, and the third instruction executes on the twelfth cycle—all of them in parallel with the cache line fill.

Each cache line fill is aligned on a 32-byte boundary. When the requested instruction or data is not 32-byte aligned, the requested item is always

loaded in the first read; each read is forwarded to the core as the line is filled. Sequential memory accesses miss the cache only when they reach the end of a cache line.

When on-chip L2 memory is configured as non-cacheable, instruction fetches and data fetches occur in 64-bit fills. In this case, each fill takes seven core cycles to complete. As shown in Figure 6-17, on-chip L2 memory is configured as non-cacheable. To illustrate the concept of L2 latency with cache off, simple instructions are used that do not require additional external data fetches. In this case, consecutive instructions are issued on consecutive cycles if multiple instructions are brought into the core in a given fetch.

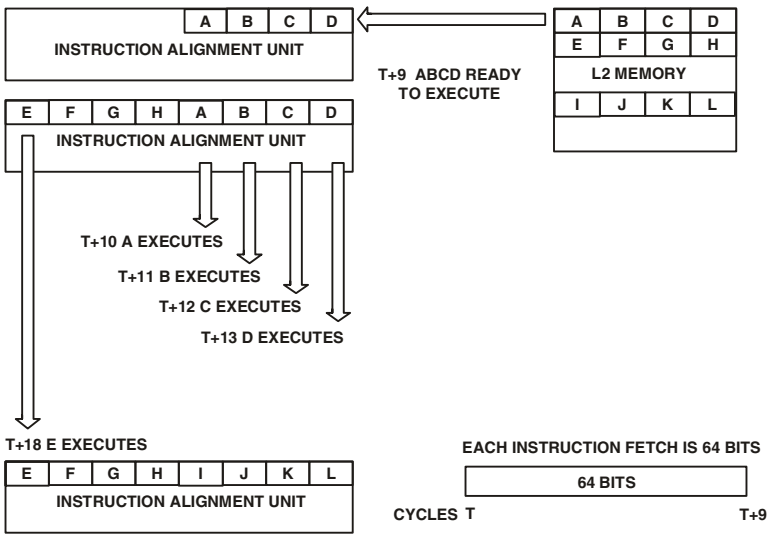

Figure 6-17. L2 Latency With Cache Off

External Memory

The external memory space is shown in Figure 6-1 on page 6-3. One of the memory regions is dedicated to four banks of SDRAM support. The size of each SDRAM bank is programmable and can range in size from 16MB to 128MB. The start address of the bank is 0x0000 0000.

Each of the next four banks contains 64MB and is dedicated to support asynchronous memories. The start address of the asynchronous memory bank is 0x2000 0000.

Memory Protection and Properties

This section describes the Memory Management Unit (MMU), memory pages, CPLB management, MMU management, and CPLB registers.

Memory Management Unit

The Blackfin processor contains a page based Memory Management Unit (MMU). This mechanism provides control over cacheability of memory ranges, as well as management of protection attributes at a page level. The MMU provides great flexibility in allocating memory and I/O resources between tasks, with complete control over access rights and cache behavior.

The MMU is implemented as two 16-entry Content Addressable Memory (CAM) blocks. Each entry is referred to as a Cacheability Protection Lookaside Buffer (CPLB) descriptor. When enabled, every valid entry in the MMU is examined on any fetch, load, or store operation to determine whether there is a match between the address being requested and the page described by the CPLB entry. If a match occurs, the cacheability and protection attributes contained in the descriptor are used for the memory transaction with no additional cycles added to the execution of the instruction.

Because the L1 memories are separated into instruction and data memories, the CPLB entries are also divided between instruction and data CPLBs. Sixteen CPLB entries are used for instruction fetch requests; these are called *ICPLBs*. Another sixteen CPLB entries are used for data transactions; these are called *DCPLBs*. The ICPLBs and DCPLBs are enabled by setting the appropriate bits in the L1 Instruction Memory Control (`IMEM_CONTROL`) and L1 Data Memory Control (`DMEM_CONTROL`) registers, respectively. These registers are shown in Figure 6-2 on page 6-8 and Figure 6-9 on page 6-27.

Each CPLB entry consists of a pair of 32-bit values. For instruction fetches:

- `ICPLB_ADDR[n]` defines the start address of the page described by the CPLB descriptor.
- `ICPLB_DATA[n]` defines the properties of the page described by the CPLB descriptor.

For data operations:

- `DCPLB_ADDR[m]` defines the start address of the page described by the CPLB descriptor.
- `DCPLB_DATA[m]` defines the properties of the page described by the CPLB descriptor.

There are two default CPLB descriptors for data accesses to the scratchpad data memory and to the system and core MMR space. These default descriptors define the above space as non-cacheable, so that additional CPLBs do not need to be set up for these regions of memory.

If valid CPLBs are set up for this space, the default CPLBs are ignored.

Memory Pages

The 4 GB address space of the processor can be divided into smaller ranges of memory or I/O referred to as memory pages. Every address within a page shares the attributes defined for that page. The architecture supports four different page sizes:

- 1 KB
- 4 KB
- 1 MB
- 4 MB

Different page sizes provide a flexible mechanism for matching the mapping of attributes to different kinds of memory and I/O.

Memory Page Attributes

Each page is defined by a two-word descriptor, consisting of an address descriptor word `xCPLB_ADDR[n]` and a properties descriptor word `xCPLB_DATA[n]`. The address descriptor word provides the base address of the page in memory. Pages must be aligned on page boundaries that are an integer multiple of their size. For example, a 4 MB page must start on an address divisible by 4 MB; whereas a 1 KB page can start on any 1 KB boundary. The second word in the descriptor specifies the other properties or attributes of the page. These properties include:

- Page size

  1 KB, 4 KB, 1 MB, 4 MB

- Cacheable/non-cacheable

  Accesses to this page use the L1 cache or bypass the cache.

- If cacheable: write-through/write-back

  Data writes propagate directly to memory or are deferred until the cache line is reallocated. If write-through, allocate on read only, or read and write.

- Dirty/modified

  The data in this page in memory has changed since the CPLB was last loaded.

- Supervisor write access permission

  – Enables or disables writes to this page when in Supervisor mode.

  – Data pages only.

- User write access permission

  – Enables or disables writes to this page when in User mode.

  – Data pages only.

- User read access permission

  Enables or disables reads from this page when in User mode.

- Valid

  Check this bit to determine whether this is valid CPLB data.

- Lock

  Keep this entry in MMR; do not participate in CPLB replacement policy.

Page Descriptor Table

For memory accesses to utilize the cache when CPLBs are enabled for instruction access, data access, or both, a valid CPLB entry must be

available in an MMR pair. The MMR storage locations for CPLB entries are limited to 16 descriptors for instruction fetches and 16 descriptors for data load and store operations.

For small and/or simple memory models, it may be possible to define a set of CPLB descriptors that fit into these 32 entries, cover the entire addressable space, and never need to be replaced. This type of definition is referred to as a *static* memory management model.

However, operating environments commonly define more CPLB descriptors to cover the addressable memory and I/O spaces than will fit into the available on-chip CPLB MMRs. When this happens, a memory-based data structure, called a Page Descriptor Table, is used; in it can be stored all the potentially required CPLB descriptors. The specific format for the Page Descriptor Table is not defined as part of the Blackfin processor architecture. Different operating systems, which have different memory management models, can implement Page Descriptor Table structures that are consistent with the OS requirements. This allows adjustments to be made between the level of protection afforded versus the performance attributes of the memory-management support routines.

CPLB Management

When the Blackfin processor issues a memory operation for which no valid CPLB (cacheability protection look aside buffer) descriptor exists in an MMR pair, an exception occurs that places the processor into Supervisor mode and vectors to the MMU exception handler (see "Exceptions" on page 4-43 for more information). The handler is typically part of the operating system (OS) kernel that implements the CPLB replacement policy.

Before CPLBs are enabled, valid CPLB descriptors must be in place for both the Page Descriptor Table and the MMU exception handler. The `LOCK` bits of these CPLB descriptors are commonly set so they are not inadvertently replaced in software.

The handler uses the faulting address to index into the Page Descriptor Table structure to find the correct CPLB descriptor data to load into one of the on-chip CPLB register pairs. If all on-chip registers contain valid CPLB entries, the handler selects one of the descriptors to be replaced, and the new descriptor information is loaded. Before loading new descriptor data into any CPLBs, the corresponding group of sixteen CPLBs must be disabled using:

- The Enable DCPLB (ENDCPLB) bit in the DMEM_CONTROL register for data descriptors, or
- The Enable ICPLB (ENICPLB) bit in the IMEM_CONTROL register for instruction descriptors

The CPLB replacement policy and algorithm to be used are the responsibility of the system MMU exception handler. This policy, which is dictated by the characteristics of the operating system, usually implements a modified LRU (Least Recently Used) policy, a round robin scheduling method, or pseudo random replacement.

After the new CPLB descriptor is loaded, the exception handler returns, and the faulting memory operation is restarted. this operation should now find a valid CPLB descriptor for the requested address, and it should proceed normally.

A single instruction may generate an instruction fetch as well as one or two data accesses. It is possible that more than one of these memory operations references data for which there is no valid CPLB descriptor in an MMR pair. In this case, the exceptions are prioritized and serviced in this order:

- Instruction page miss
- A page miss on DAG0
- A page miss on DAG1

MMU Application

Memory management is an optional feature in the Blackfin processor architecture. Its use is predicated on the system requirements of a given application. Upon reset, all CPLBs are disabled, and the Memory Management Unit (MMU) is not used.

If all L1 memory is configured as SRAM, then the data and instruction MMU functions are optional, depending on the application's need for protection of memory spaces either between tasks or between User and Supervisor modes. To protect memory between tasks, the operating system can maintain separate tables of instruction and/or data memory pages available for each task and make those pages visible only when the relevant task is running. When a task switch occurs, the operating system can ensure the invalidation of any CPLB descriptors on chip that should not be available to the new task. It can also preload descriptors appropriate to the new task.

For many operating systems, the application program is run in User mode while the operating system and its services run in Supervisor mode. It is desirable to protect code and data structures used by the operating system from inadvertent modification by a running User mode application. This protection can be achieved by defining CPLB descriptors for protected memory ranges that allow write access only when in Supervisor mode. If a write to a protected memory region is attempted while in User mode, an exception is generated before the memory is modified. Optionally, the User mode application may be granted read access for data structures that are useful to the application. Even Supervisor mode functions can be blocked from writing some memory pages that contain code that is not expected to be modified. Because CPLB entries are MMRs that can be written only while in Supervisor mode, user programs cannot gain access to resources protected in this way.

If either the L1 Instruction Memory or the L1 Data Memory is configured partially or entirely as cache, the corresponding CPLBs must be enabled. When an instruction generates a memory request and the cache is enabled,

the processor first checks the ICPLBs to determine whether the address requested is in a cacheable address range. If no valid ICPLB entry in an MMR pair corresponds to the requested address, an MMU exception is generated to obtain a valid ICPLB descriptor to determine whether the memory is cacheable or not. As a result, if the L1 Instruction Memory is enabled as cache, then any memory region that contains instructions must have a valid ICPLB descriptor defined for it. These descriptors must either reside in MMRs at all times or be resident in a memory-based Page Descriptor Table that is managed by the MMU exception handler. Likewise, if either or both L1 data banks are configured as cache, all potential data memory ranges must be supported by DCPLB descriptors.

Before caches are enabled, the MMU and its supporting data structures must be set up and enabled.

Examples of Protected Memory Regions

In Figure 6-18, a starting point is provided for basic CPLB allocation for Instruction and Data CPLBs. Note some ICPLBs and DCPLBs have common descriptors for the same address space.

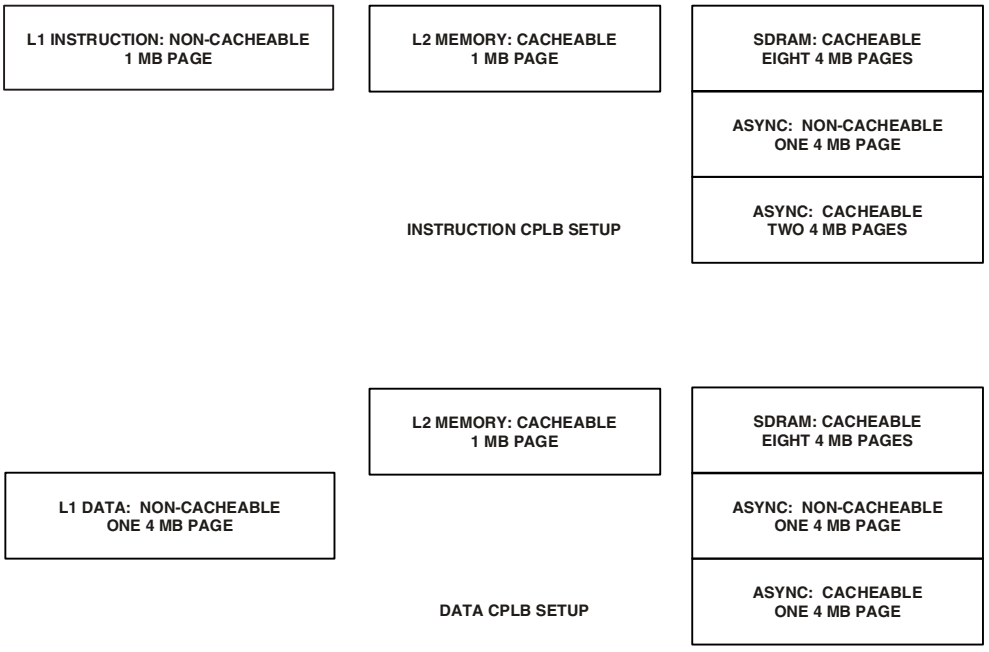

Figure 6-18. Examples of Protected Memory Regions

ICPLB Data Registers (ICPLB_DATAx)

Figure 6-19 describes the ICPLB Data registers.

To ensure proper behavior and future compatibility, all reserved bits in this register must be set to 0 whenever this register is written.

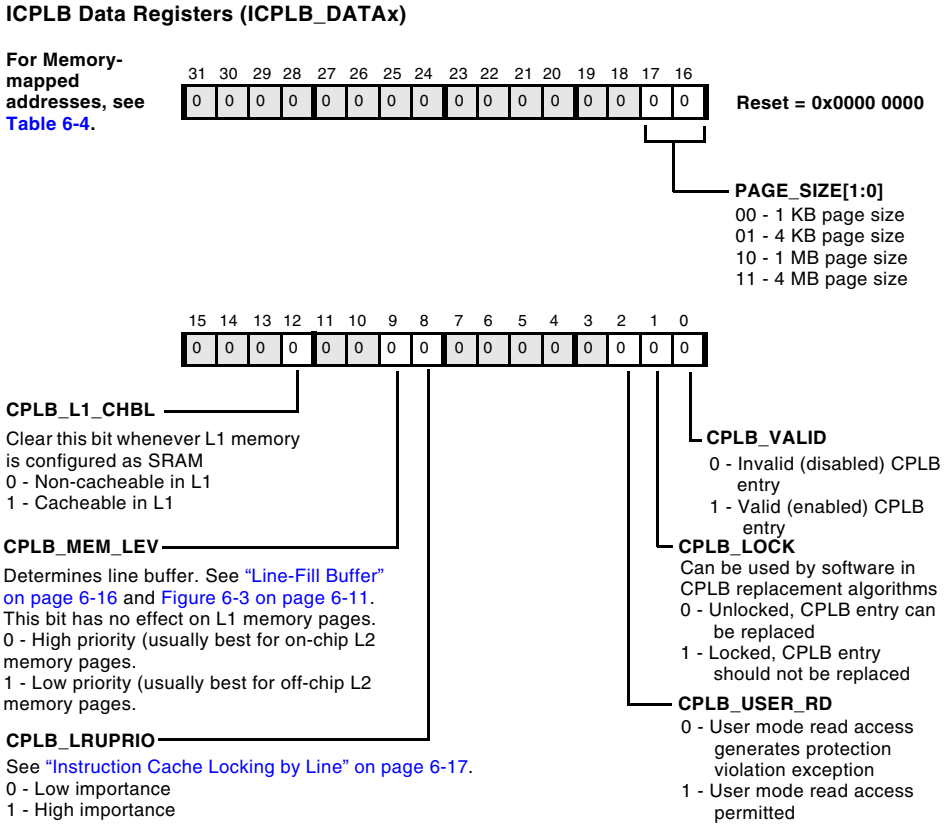

Figure 6-19. ICPLB Data Registers

Table 6-4. ICPLB Data Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
ICPLB_DATA0	0xFFE0 1200
ICPLB_DATA1	0xFFE0 1204
ICPLB_DATA2	0xFFE0 1208
ICPLB_DATA3	0xFFE0 120C
ICPLB_DATA4	0xFFE0 1210
ICPLB_DATA5	0xFFE0 1214
ICPLB_DATA6	0xFFE0 1218
ICPLB_DATA7	0xFFE0 121C
ICPLB_DATA8	0xFFE0 1220
ICPLB_DATA9	0xFFE0 1224
ICPLB_DATA10	0xFFE0 1228
ICPLB_DATA11	0xFFE0 122C
ICPLB_DATA12	0xFFE0 1230
ICPLB_DATA13	0xFFE0 1234
ICPLB_DATA14	0xFFE0 1238
ICPLB_DATA15	0xFFE0 123C

DCPLB Data Registers (DCPLB_DATAx)

Figure 6-20 shows the DCPLB Data registers.

To ensure proper behavior and future compatibility, all reserved bits in this register must be set to 0 whenever this register is written.

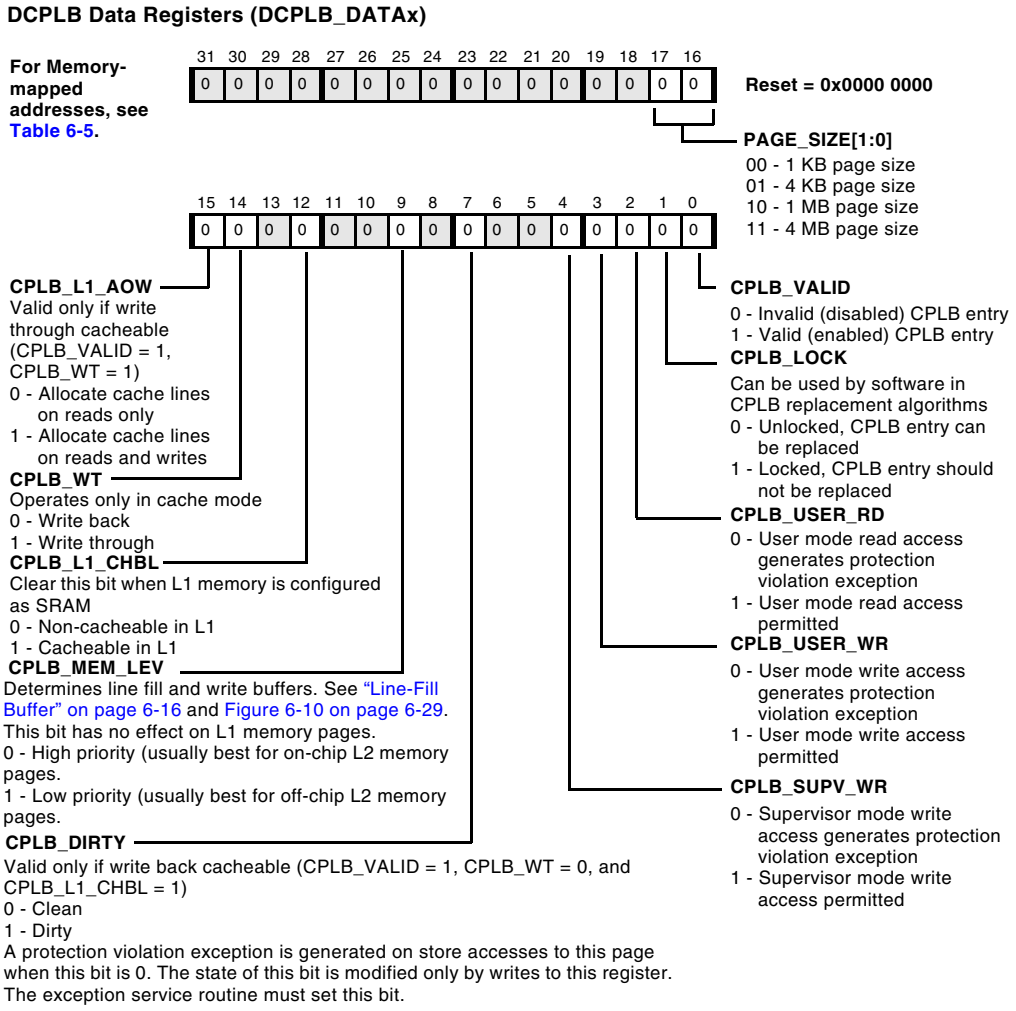

Figure 6-20. DCPLB Data Registers

Table 6-5. DCPLB Data Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
DCPLB_DATA0	0xFFE0 0200
DCPLB_DATA1	0xFFE0 0204
DCPLB_DATA2	0xFFE0 0208
DCPLB_DATA3	0xFFE0 020C
DCPLB_DATA4	0xFFE0 0210
DCPLB_DATA5	0xFFE0 0214
DCPLB_DATA6	0xFFE0 0218
DCPLB_DATA7	0xFFE0 021C
DCPLB_DATA8	0xFFE0 0220
DCPLB_DATA9	0xFFE0 0224
DCPLB_DATA10	0xFFE0 0228
DCPLB_DATA11	0xFFE0 022C
DCPLB_DATA12	0xFFE0 0230
DCPLB_DATA13	0xFFE0 0234
DCPLB_DATA14	0xFFE0 0238
DCPLB_DATA15	0xFFE0 023C

DCPLB Address Registers (DCPLB_ADDRx)

Figure 6-21 shows the DCPLB Address registers.

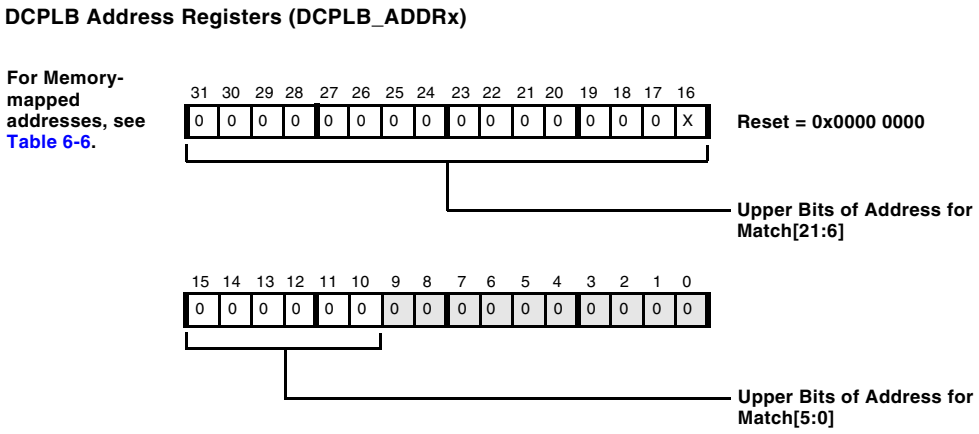

Figure 6-21. DCPLB Address Registers

Table 6-6. DCPLB Address Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
DCPLB_ADDR0	0xFFE0 0100
DCPLB_ADDR1	0xFFE0 0104
DCPLB_ADDR2	0xFFE0 0108
DCPLB_ADDR3	0xFFE0 010C
DCPLB_ADDR4	0xFFE0 0110
DCPLB_ADDR5	0xFFE0 0114
DCPLB_ADDR6	0xFFE0 0118
DCPLB_ADDR7	0xFFE0 011C
DCPLB_ADDR8	0xFFE0 0120
DCPLB_ADDR9	0xFFE0 0124

Table 6-6. DCPLB Address Register Memory-Mapped Addresses (Cont'd)

Register Name	Memory-Mapped Address
DCPLB_ADDR10	0xFFE0 0128
DCPLB_ADDR11	0xFFE0 012C
DCPLB_ADDR12	0xFFE0 0130
DCPLB_ADDR13	0xFFE0 0134
DCPLB_ADDR14	0xFFE0 0138
DCPLB_ADDR15	0xFFE0 013C

ICPLB Address Registers (ICPLB_ADDRx)

Figure 6-22 shows the ICPLB Address registers.

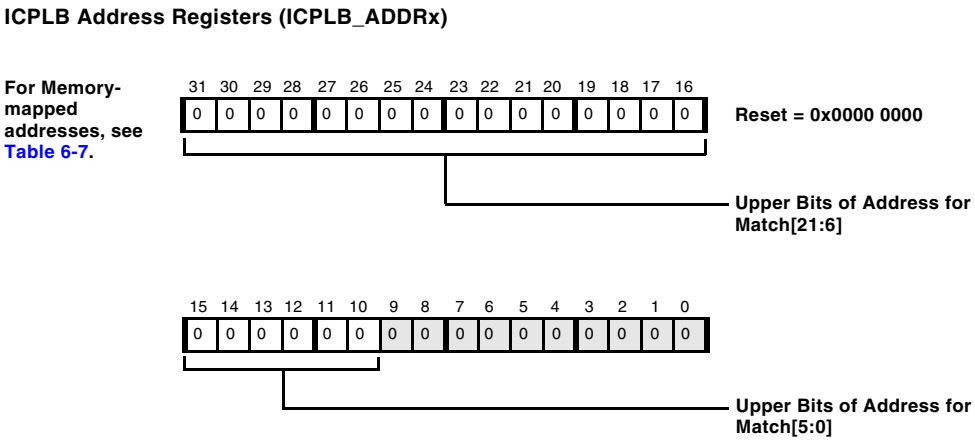

Figure 6-22. ICPLB Address Registers

Table 6-7. ICPLB Address Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
ICPLB_ADDR0	0xFFE0 1100
ICPLB_ADDR1	0xFFE0 1104
ICPLB_ADDR2	0xFFE0 1108
ICPLB_ADDR3	0xFFE0 110C
ICPLB_ADDR4	0xFFE0 1110
ICPLB_ADDR5	0xFFE0 1114
ICPLB_ADDR6	0xFFE0 1118
ICPLB_ADDR7	0xFFE0 111C
ICPLB_ADDR8	0xFFE0 1120
ICPLB_ADDR9	0xFFE0 1124
ICPLB_ADDR10	0xFFE0 1128
ICPLB_ADDR11	0xFFE0 112C
ICPLB_ADDR12	0xFFE0 1130
ICPLB_ADDR13	0xFFE0 1134
ICPLB_ADDR14	0xFFE0 1138
ICPLB_ADDR15	0xFFE0 113C

DCPLB and ICPLB Status Registers (DCPLB_STATUS, ICPLB_STATUS)

Bits in the DCPLB Status register (`DCPLB_STATUS`) and ICPLB Status register (`ICPLB_STATUS`) identify the CPLB entry that has triggered CPLB-related exceptions. The exception service routine can infer the cause of the fault by examining the CPLB entries.

The `DCPLB_STATUS` and `ICPLB_STATUS` registers are valid only while in the faulting exception service routine.

Bits FAULT_DAG, FAULT_USERSUPV and FAULT_RW in the DCPLB Status register (DCPLB_STATUS) are used to identify the CPLB entry that has triggered the CPLB-related exception (see Figure 6-23).

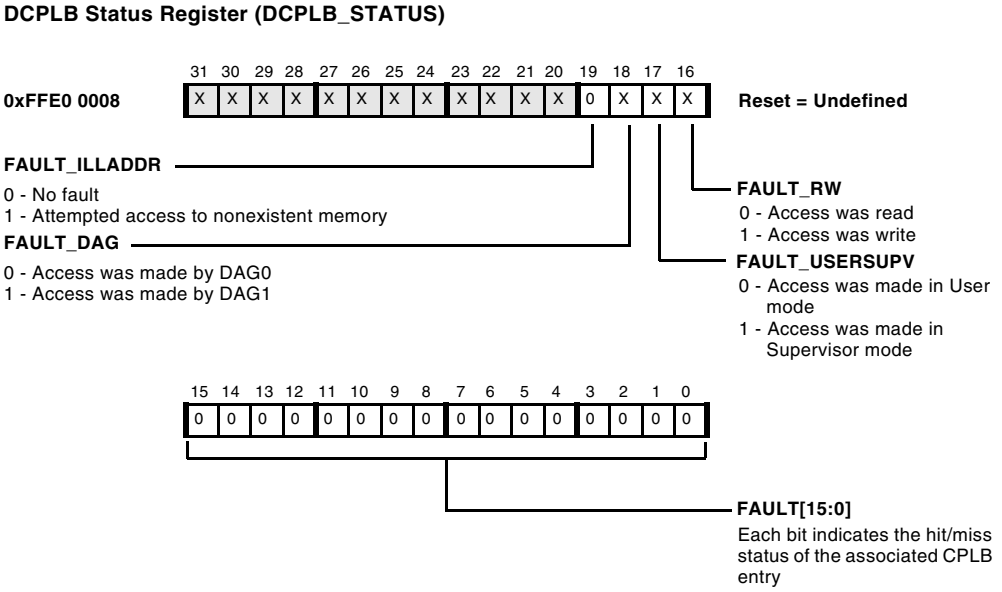

Figure 6-23. DCPLB Status Register

Bit FAULT_USERSUPV in the ICPLB Status register (ICPLB_STATUS) is used to identify the CPLB entry that has triggered the CPLB-related exception (see Figure 6-24).

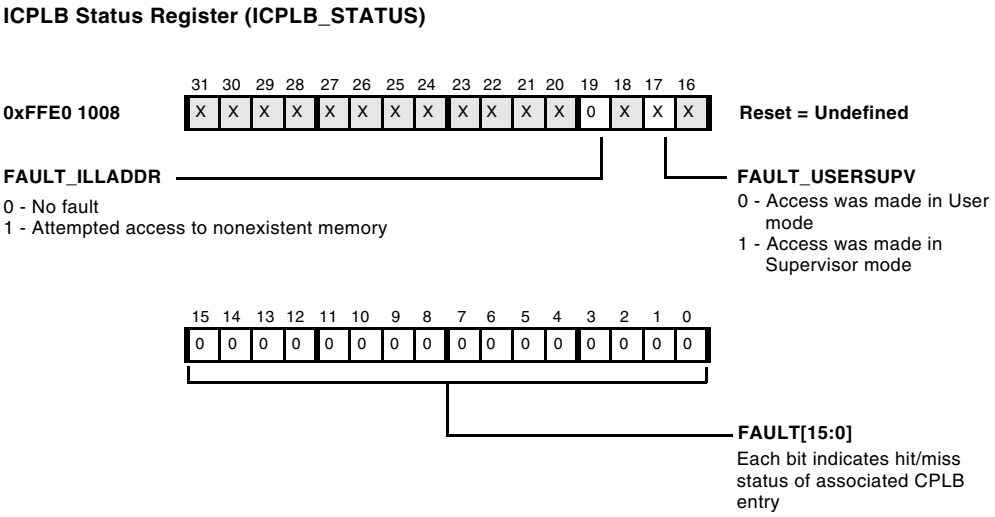

Figure 6-24. ICPLB Status Register

DCPLB and ICPLB Fault Address Registers (DCPLB_FAULT_ADDR, ICPLB_FAULT_ADDR)

The DCPLB Address register (DCPLB_FAULT_ADDR) and ICPLB Fault Address register (ICPLB_FAULT_ADDR) hold the address that has caused a fault in the L1 Data Memory or L1 Instruction Memory, respectively. See Figure 6-25 and Figure 6-26.

The DCPLB_FAULT_ADDR and ICPLB_FAULT_ADDR registers are valid only while in the faulting exception service routine.

DCPLB Address Register (DCPLB_FAULT_ADDR)

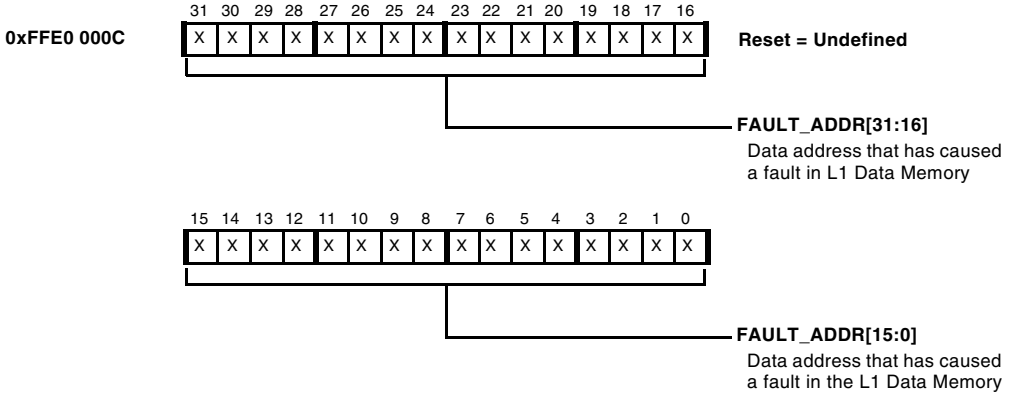

Figure 6-25. DCPLB Address Register

ICPLB Fault Address Register (ICPLB_FAULT_ADDR)

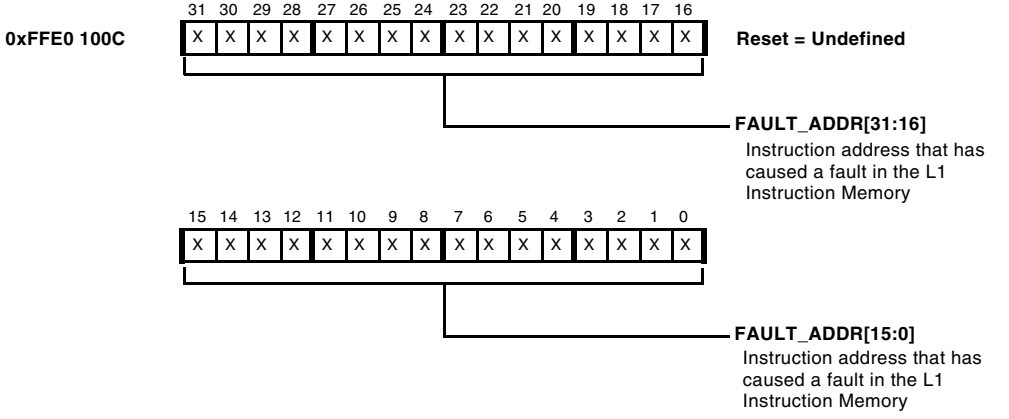

Figure 6-26. ICPLB Fault Address Register

Memory Transaction Model

Both internal and external memory locations are accessed in little endian byte order. Figure 6-27 shows a data word stored in register R0 and in memory at address location *addr*. B0 refers to the least significant byte of the 32-bit word.

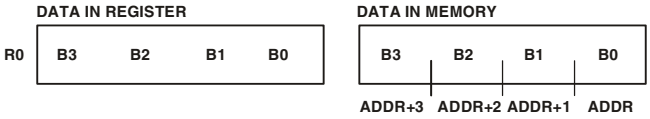

Figure 6-27. Data Stored in Little Endian Order

Figure 6-28 shows 16- and 32-bit instructions stored in memory. The diagram on the left shows 16-bit instructions stored in memory with the most significant byte of the instruction stored in the high address (byte B1 in *addr+1*) and the least significant byte in the low address (byte B0 in *addr*).

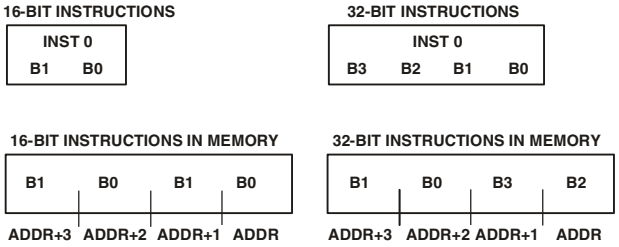

Figure 6-28. Instructions Stored in Little Endian Order

The diagram on the right shows 32-bit instructions stored in memory. Note the most significant 16-bit half word of the instruction (bytes B3 and B2) is stored in the low addresses (*addr+1* and *addr*), and the least significant half word (bytes B1 and B0) is stored in the high addresses (*addr+3* and *addr+2*).

Load/Store Operation

The Blackfin processor architecture supports the RISC concept of a Load/Store machine. This machine is the characteristic in RISC architectures whereby memory operations (loads and stores) are intentionally separated from the arithmetic functions that use the targets of the memory operations. The separation is made because memory operations, particularly instructions that access off-chip memory or I/O devices, often take multiple cycles to complete and would normally halt the processor, preventing an instruction execution rate of one instruction per cycle.

Separating load operations from their associated arithmetic functions allows compilers or assembly language programmers to place unrelated instructions between the load and its dependent instructions. The unrelated instructions execute in parallel while the processor waits for the memory system to return the data. If the value is returned before the dependent operation reaches the execution stage of the pipeline, the operation completes in one cycle.

In write operations, the store instruction is considered complete as soon as it executes, even though many cycles may execute before the data is actually written to an external memory or I/O location. This arrangement allows the processor to execute one instruction per clock cycle, and it implies that the synchronization between when writes complete and when subsequent instructions execute is not guaranteed. Moreover, this synchronization is considered unimportant in the context of most memory operations.

Interlocked Pipeline

In the execution of instructions, the Blackfin processor architecture implements an interlocked pipeline. When a load instruction executes, the target register of the read operation is marked as busy until the value is returned from the memory system. If a subsequent instruction tries to access this register before the new value is present, the pipeline will stall

until the memory operation completes. This stall guarantees that instructions that require the use of data resulting from the load do not use the previous or invalid data in the register, even though instructions are allowed to start execution before the memory read completes.

This mechanism allows the execution of independent instructions between the load and the instructions that use the read target without requiring the programmer or compiler to know how many cycles are actually needed for the memory-read operation to complete. If the instruction immediately following the load uses the same register, it simply stalls until the value is returned. Consequently, it operates as the programmer expects. However, if four other instructions are placed after the load but before the instruction that uses the same register, all of them execute, and the overall throughput of the processor is improved.

Ordering of Loads and Stores

The relaxation of synchronization between memory access instructions and their surrounding instructions is referred to as weak ordering of loads and stores. Weak ordering implies that the timing of the actual completion of the memory operations—even the order in which these events occur—may not align with how they appear in the sequence of the program source code. All that is guaranteed is:

- Load operations will complete before the returned data is used by a subsequent instruction.
- Load operations using data previously written will use the updated values.
- Store operations will eventually propagate to their ultimate destination.

Because of weak ordering, the memory system is allowed to prioritize reads over writes. In this case, a write that is queued anywhere in the pipeline, but not completed, may be deferred by a subsequent read operation,

and the read is allowed to be completed before the write. Reads are prioritized over writes because the read operation has a dependent operation waiting on its completion, whereas the processor considers the write operation complete, and the write does not stall the pipeline if it takes more cycles to propagate the value out to memory. This behavior could cause a read that occurs in the program source code after a write in the program flow to actually return its value before the write has been completed. This ordering provides significant performance advantages in the operation of most memory instructions. However, it can cause side effects that the programmer must be aware of to avoid improper system operation.

When writing to or reading from non memory locations such as I/O device registers, the order of how read and write operations complete is often significant. For example, a read of a status register may depend on a write to a control register. If the address is the same, the read would return a value from the write buffer rather than from the actual I/O device register, and the order of the read and write at the register may be reversed. Both these effects could cause undesirable side effects in the intended operation of the program and peripheral. To ensure that these effects do not occur in code that requires precise (strong) ordering of load and store operations, synchronization instructions (`CSYNC` or `SSYNC`) should be used.

Synchronizing Instructions

When strong ordering of loads and stores is required, as may be the case for sequential writes to an I/O device for setup and control, use the core or system synchronization instructions, `CSYNC` or `SSYNC`, respectively.

The `CSYNC` instruction ensures all pending core operations have completed and the core buffer (between the processor core and the L1 memories) has been flushed before proceeding to the next instruction. Pending core operations may include any pending interrupts, speculative states (such as branch predictions), or exceptions.

Consider the following example code sequence:

```
IF CC JUMP away_from_here
csync;
r0 = [p0];
away_from_here:
```

In the preceding example code, the CSYNC instruction ensures:

- The conditional branch (IF CC JUMP away_from_here) is resolved, forcing stalls into the execution pipeline until the condition is resolved and any entries in the processor store buffer have been flushed.
- All pending interrupts or exceptions have been processed before CSYNC completes.
- The load is not fetched from memory speculatively.

The SSYNC instruction ensures that all side effects of previous operations are propagated out through the interface between the L1 memories and the rest of the chip. In addition to performing the core synchronization functions of CSYNC, the SSYNC instruction flushes any write buffers between the L1 memory and the system domain and generates a sync request to the system that requires acknowledgement before SSYNC completes.

Speculative Load Execution

Load operations from memory do not change the state of the memory value. Consequently, issuing a speculative memory-read operation for a subsequent load instruction usually has no undesirable side effect. In some code sequences, such as a conditional branch instruction followed by a load, performance may be improved by speculatively issuing the read

request to the memory system before the conditional branch is resolved. For example,

```
            IF CC JUMP away_from_here
            R0 = [P2];
            …
away_from_here:
```

If the branch is taken, then the load is flushed from the pipeline, and any results that are in the process of being returned can be ignored. Conversely, if the branch is not taken, the memory will have returned the correct value earlier than if the operation were stalled until the branch condition was resolved.

However, in the case of an I/O device, this could cause an undesirable side effect for a peripheral that returns sequential data from a FIFO or from a register that changes value based on the number of reads that are requested. To avoid this effect, use synchronizing instructions (`CSYNC` or `SSYNC`) to guarantee the correct behavior between read operations.

Store operations never access memory speculatively, because this could cause modification of a memory value before it is determined whether the instruction should have executed.

Conditional Load Behavior

The synchronization instructions force all speculative states to be resolved before a load instruction initiates a memory reference. However, the load instruction itself may generate more than one memory-read operation, because it is interruptible. If an interrupt of sufficient priority occurs between the completion of the synchronization instruction and the completion of the load instruction, the sequencer cancels the load instruction. After execution of the interrupt, the interrupted load is executed again. This approach minimizes interrupt latency. However, it is possible that a memory-read cycle was initiated before the load was canceled, and this would be followed by a second read operation after the load is executed

again. For most memory accesses, multiple reads of the same memory address have no side effects. However, for some memory-mapped devices, such as peripheral data FIFOs, reads are destructive. Each time the device is read, the FIFO advances, and the data cannot be recovered and re-read.

When accessing memory-mapped devices that have state dependencies on the number of read operations on a given address location, disable interrupts before performing the load operation.

Working With Memory

This section contains information about alignment of data in memory and memory operations that support semaphores between tasks. It also contains a brief discussion of MMR registers and a core MMR programming example.

Alignment

Nonaligned memory operations are not directly supported. A nonaligned memory reference generates a Misaligned Access exception event (see "Exceptions" on page 4-43). However, because some data streams (such as 8-bit video data) can properly be nonaligned in memory, alignment exceptions may be disabled by using the `DISALGNEXCPT` instruction. Moreover, some instructions in the quad 8-bit group automatically disable alignment exceptions.

Cache Coherency

For shared data, software must provide cache coherency support as required. To accomplish this, use the `FLUSH` instruction (see "Data Cache Control Instructions" on page 6-38), and/or explicit line invalidation through the core MMRs (see "Data Test Registers" on page 6-39).

Atomic Operations

The processor provides a single atomic operation: TESTSET. Atomic operations are used to provide non interruptible memory operations in support of semaphores between tasks. The TESTSET instruction loads an indirectly addressed memory half word, tests whether the low byte is zero, and then sets the most significant bit (MSB) of the low memory byte without affecting any other bits. If the byte is originally zero, the instruction sets the CC bit. If the byte is originally nonzero, the instruction clears the CC bit. The sequence of this memory transaction is atomic—hardware bus locking insures that no other memory operation can occur between the test and set portions of this instruction. The TESTSET instruction can be interrupted by the core. If this happens, the TESTSET instruction is executed again upon return from the interrupt.

The TESTSET instruction can address the entire 4 GB memory space, but should not target on-core memory (L1 or MMR space) since atomic access to this memory is not supported.

The memory architecture always treats atomic operations as cache inhibited accesses even if the CPLB descriptor for the address indicates cache enabled access. However, executing TESTSET operations on cacheable regions of memory is not recommended since the architecture cannot guarantee a cacheable location of memory is coherent when the TESTSET instruction is executed.

Memory-Mapped Registers

The MMR reserved space is located at the top of the memory space (0xFFC0 0000). This region is defined as non-cacheable and is divided between the system MMRs (0xFFC0 0000–0xFFE0 0000) and core MMRs (0xFFE0 0000–0xFFFF FFFF).

If strong ordering is required, place a synchronization instruction after stores to MMRs. For more information, see "Load/Store Operation" on page 6-67.

All MMRs are accessible only in Supervisor mode. Access to MMRs in User mode generates a protection violation exception. Attempts to access MMR space using DAG1 will generate a protection violation exception.

All core MMRs are read and written using 32-bit aligned accesses. However, some MMRs have fewer than 32 bits defined. In this case, the unused bits are reserved. System MMRs may be 16 bits.

Accesses to nonexistent MMRs generate an illegal access exception. The system ignores writes to read-only MMRs.

Appendix A, "Blackfin Processor Core MMR Assignments", provides a summary of all Core MMRs. Appendix B, "System MMR Assignments", provides a summary of all System MMRs.

Core MMR Programming Code Example

Core MMRs may be accessed only as aligned 32-bit words. Nonaligned access to MMRs generates an exception event. Listing 6-1 shows the instructions required to manipulate a generic core MMR.

Listing 6-1. Core MMR Programming

```
CLI R0;   /* stop interrupts and save IMASK */
P0 = MMR_BASE;   /* 32-bit instruction to load base of MMRs */
R1 = [P0 + TIMER_CONTROL_REG];   /* get value of control reg */
BITSET R1, #N;   /* set bit N */
[P0 + TIMER_CONTROL_REG] = R1;   /* restore control reg */
```

```
CSYNC;   /* assures that the control reg is written */
STI R0;   /* enable interrupts */
```

The CLI instruction saves the contents of the IMASK register and disables interrupts by clearing IMASK. The STI instruction restores the contents of the IMASK register, thus enabling interrupts. The instructions between CLI and STI are not interruptible.

Terminology

The following terminology is used to describe memory.

cache block. The smallest unit of memory that is transferred to/from the next level of memory from/to a cache as a result of a cache miss.

cache hit. A memory access that is satisfied by a valid, present entry in the cache.

cache line. Same as cache block. In this chapter, cache line is used for cache block.

cache miss. A memory access that does not match any valid entry in the cache.

direct-mapped. Cache architecture in which each line has only one place in which it can appear in the cache. Also described as 1-Way associative.

dirty or **modified**. A state bit, stored along with the tag, indicating whether the data in the data cache line has been changed since it was copied from the source memory and, therefore, needs to be updated in that source memory.

exclusive, clean. The state of a data cache line, indicating that the line is valid and that the data contained in the line matches that in the source memory. The data in a clean cache line does not need to be written to source memory before it is replaced.

fully associative. Cache architecture in which each line can be placed anywhere in the cache.

index. Address portion that is used to select an array element (for example, a line index).

invalid. Describes the state of a cache line. When a cache line is invalid, a cache line match cannot occur.

least recently used (LRU) algorithm. Replacement algorithm, used by cache, that first replaces lines that have been unused for the longest time.

Level 1 (L1) memory. Memory that is directly accessed by the core with no intervening memory subsystems between it and the core.

little endian. The native data store format of the Blackfin processor. Words and half words are stored in memory (and registers) with the least significant byte at the lowest byte address and the most significant byte in the highest byte address of the data storage location.

replacement policy. The function used by the processor to determine which line to replace on a cache miss. Often, an LRU algorithm is employed.

set. A group of *N*-line storage locations in the Ways of an *N*-Way cache, selected by the `INDEX` field of the address (see Figure 6-4 on page 6-13).

set associative. Cache architecture that limits line placement to a number of sets (or Ways).

tag. Upper address bits, stored along with the cached data line, to identify the specific address source in memory that the cached line represents.

valid. A state bit, stored with the tag, indicating that the corresponding tag and data are current and correct and can be used to satisfy memory access requests.

victim. A dirty cache line that must be written to memory before it can be replaced to free space for a cache line allocation.

Way. An array of line storage elements in an *N*-Way cache (see Figure 6-4 on page 6-13).

write back. A cache write policy, also known as *copyback*. The write data is written only to the cache line. The modified cache line is written to source memory only when it is replaced. Cache lines are allocated on both reads and writes.

write through. A cache write policy (also known as store through). The write data is written to both the cache line and to the source memory. The modified cache line is *not* written to the source memory when it is replaced. Cache lines must be allocated on reads, and may be allocated on writes (depending on mode).

7 CHIP BUS HIERARCHY

This chapter discusses on-chip buses, how data moves through the system and factors that determine the system organization. The chapter describes the system internal chip interfaces and discusses the system interconnects, including the interfaces between core buses and system buses.

Internal Interfaces

Figure 7-4 on page 7-9 shows the two processor cores (Core A and Core B), system boundaries, and the interfaces between them. Each core (A or B) has on-core L1 memory, including 32K Bytes of L1 SRAM instruction memory, 64K Bytes of L1 SRAM data memory and 4K Bytes of system scratch pad SRAM memory. As part of on-chip memory, there is another off-core 128K Bytes of L2 memory. External memory, such as SDRAM and Flash, can be accessed through the External Bus Interface Unit (EBIU). In the ADSP-BF561, there are nine peripherals divided into two groups, Group 1 and Group 2. Three Direct Memory Access (DMA) controllers—peripheral Direct Memory Access Controllers (DMA1, DMA2) and the Internal Memory Direct Memory Access (IMDMA) controller—support the data transfers within internal memories or between peripherals and memories. See Figure 7-1.

Each Core (A or B) has three ports connected to system and L2 memory:

- 16-bit Core (A or B) MMR Port

  This is the system Memory Mapped Register (MMR) access port. Through this port, the 16-bit Peripheral Access Bus (PAB) bus connects all off-core peripherals to either the Core A or Core B system MMR registers. Interface arbitration logic allows both cores to request access to the PAB bus at the same time.

- 64-bit Core (A or B) P Port

  This is the Core (A or B) L1 memory access port. The P Port provides the interface to the External Bus Interface Unit (EBIU) and to the L2 memory through the 32-bit External Access Bus (EAB Bus) and the 64-bit Core (A or B) L2 bus separately. The memory access request commands from the core are pipelined; no arbitration logic is needed in this interface.

- 32-bit Core (A or B) D Port

  DMA controllers (DMA1, DMA2 and IMDMA) transfer data to or from core L1 memory through this port. Since there are three DMA controllers (DMA1, DMA2 and IMDMA) that can simultaneously request access to L1 memory through this 32-bit D Port, interface arbitration logic is provided and described later in this chapter.

L2 memory also has two ports connected to the following two core buses, which run at core clock frequency (CCLK domain):

- 64-bit Core L2 bus

  This bus supports the L2 memory data/instruction accesses requested by Core A or B, as well as IMDMA data transfers within L1/L2 memory systems.

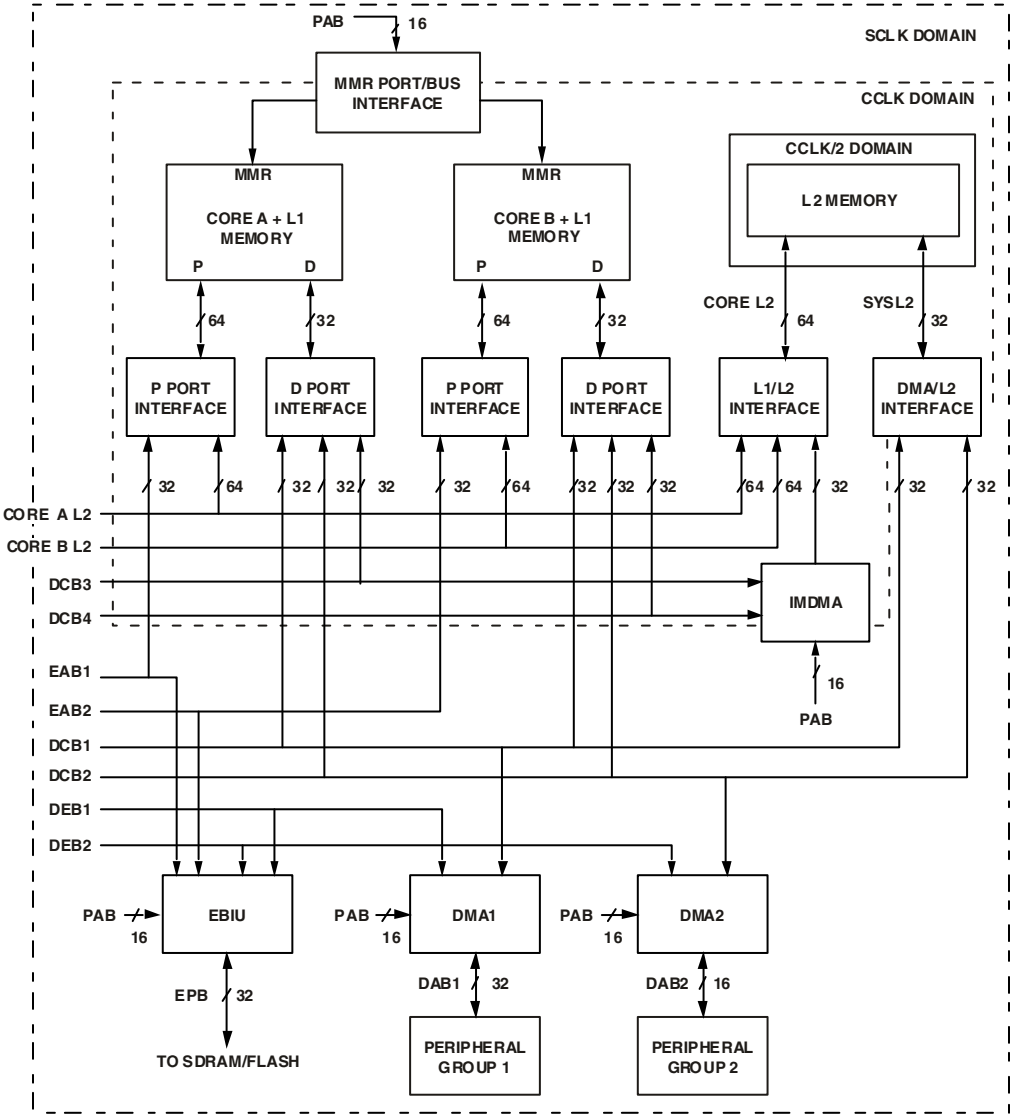

Figure 7-1. ADSP-BF561 Bus Hierarchy

- 32-bit System L2 bus (Sys L2 bus)

  This bus supports DMA1 and DMA2 data transfers between peripherals/external memory and L2 memory.

Overall system functions of the ADSP-BF561 are supported by the following system buses, which run at system clock frequency (SCLK domain):

- 16-bit Peripheral Access Bus (PAB bus)
- 32-bit External Access Bus (Core A EAB bus and Core B EAB bus)
- 32-bit DMA Core Bus (DCB1 bus, DCB2 bus, DCB3 bus and DCB4 bus)
- 32/16-bit (32-bit DAB1 bus and 16-bit DAB2 bus)
- 32-bit DMA External Bus (DEB1 bus and DEB2 bus)

The Internal Memory DMA controller (IMDMA) has two 32-bit DCB buses connected to Core A (DCB3) and Core B (DCB4) that support data transfers within internal memories L1 and L2. The DCB3 and DCB4 buses run at core clock frequency (CCLK domain). The External Port Bus (EPB bus) provides the connection between the External Bus Interface Unit (EBIU) and external memory. The EPB bus runs at system clock frequency (SCLK domain).

Peripheral Group 1 includes the following peripherals:

- PPI0
- PPI1

Peripheral Group 2 includes the following peripherals:

- SPI
- UART

- SPORT0
- SPORT1

Internal Clocks

The core processor clock (CCLK) rate is programmable with respect to CLKIN. The CCLK rate is divided down from the PLL output. The divider ratio is set using the CSEL parameter of the PLL Divide register. All the core buses, such as the Core (A or B) MMR bus, Core (A or B) I bus, Core (A or B) D0 bus, Core (A or B) D1 bus, and Core (A or B) DMA bus, run at the core clock frequency (CCLK domain). The other core buses, Core L2 bus and System L2 bus (Sys L2 bus), run at half of the core clock frequency (CCLK/2 domain). All system buses and the EBIU run at the system clock frequency (SCLK domain). The SCLK divider ratio is set using the SSEL parameter of the PLL Divide register, and must be set so that the buses run at a frequency slower or equal to the core clock frequency, and never exceeding 133 MHz. These buses can also be cycled at a programmable frequency to reduce power consumption, or to allow the core processor to run at an optimal frequency. Note that all synchronous peripherals derive their timing from SCLK.

Core Overview

Figure 7-2 shows a processor core block diagram that includes a processor core and L1 memory connected by internal core buses.

Core Interfaces

Each processor consists of a processor core and L1 memory. Core buses provide the path between the processor and L1 memory. The core bus structure between the processor core (A or B) and L1 memory runs at the

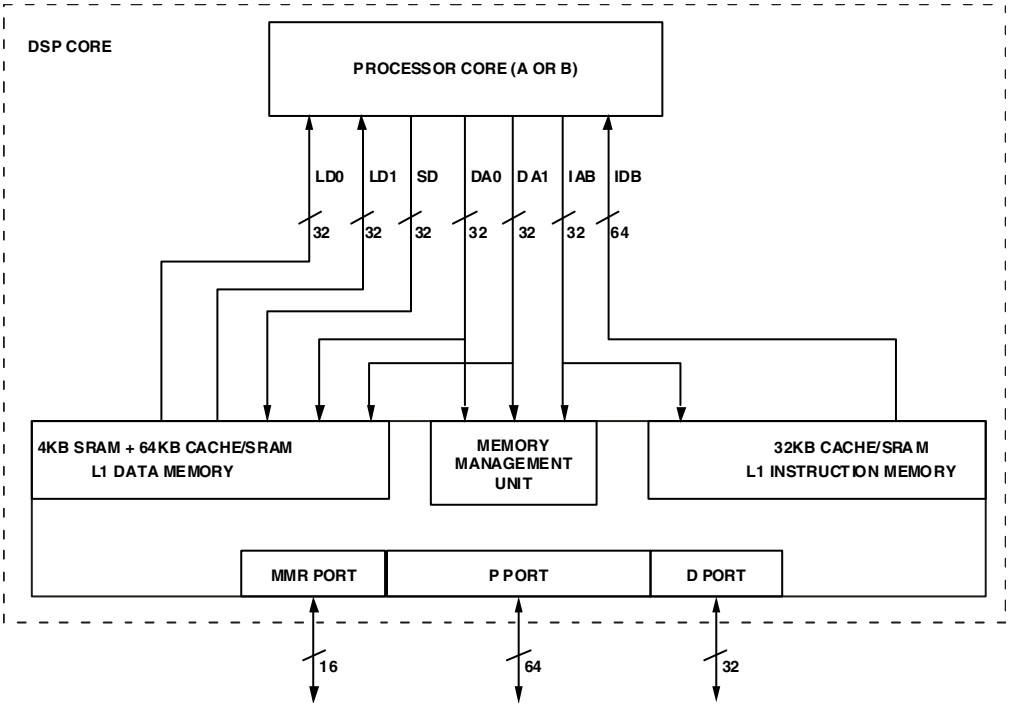

Figure 7-2. Core Block Diagram

full core frequency (CCLK domain) and allows the processor core (A or B) to perform the following L1 memory accesses per core clock cycle (CCLK):

- one 64-bit instruction fetch through the IDB bus
- one 32-bit data reference through the DA0 bus
- one 32-bit data reference through the DA1 bus
- two 32-bit data loads through the LD0 and LD1 buses
- one 32-bit data store through the SD bus

Each core (A or B) has three ports and can generate up to the following simultaneous off-core accesses per core clock cycle (CCLK):

- one DMA data transfer through the D Port
- one L2 or external memory access through the P Port
- one MMR register access through the MMR Port

The L2 or external memory access through the P port includes normal data/instruction access and cache read/write operation.

System Overview

The ADSP-BF561 system includes two Blackfin processor cores (Core A and Core B), a 128KB Level 2 (L2) memory, the peripheral set (Timers, UART, SPORTs, PPIs, Watchdog Timers and SPI), the external memory controller (EBIU AMC and SDC), the DMA controllers, and bus interfaces.

The External Bus Interface Unit (EBIU) is the primary bus to the chip pins. Detailed information about the External Bus Interface Unit is discussed in "External Bus Interface Unit" on page 16-1.

P Port Interface

Figure 7-3 shows the interface between Core (A or B) P port and L2 memory through the 64-bit Core (A or B) L2 bus; and the interface between Core (A or B) P Port and the EBIU through the 32-bit EAB bus. At each core clock cycle CCLK, the processor Core (A or B) can:

- Transfer one 64-bit instruction word from L2 memory
- Or transfer one 32/16/8-bit data word to or from the same (or different) L2 memory data bank

- Or transfer one 32/16/8-bit data word to or from external memory.

Data transfers requested from Core A or B to L2 or the EBIU are fully pipelined.

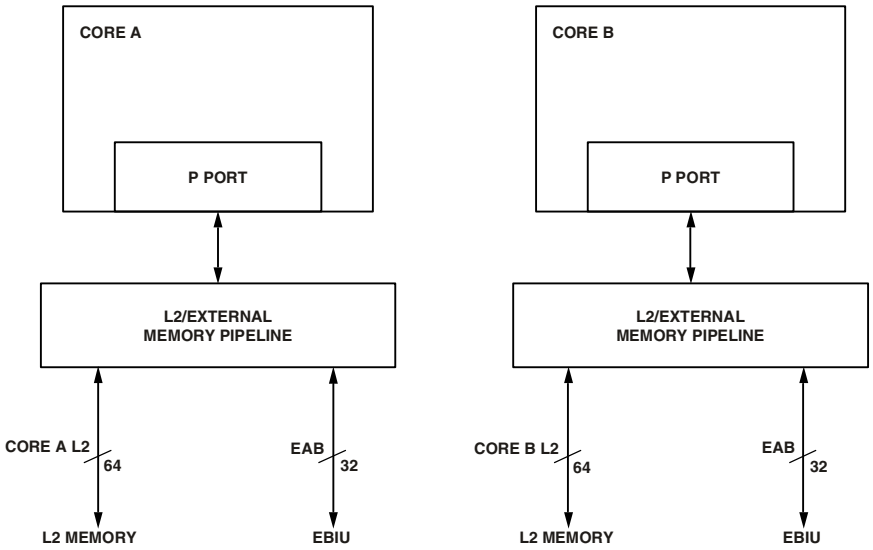

Figure 7-3. Core (A or B) Interface to L2 Memory and the EBIU

D Port Interface

Figure 7-4 shows the interface between the three DMA controller core access buses (32-bit DCB buses) and the Core's D port. This 32-bit D Port provides DMA access to L1 memory.

Three DCB buses provide the following DMA data transfers:

- The DCB1 bus supports DMA data transfer between Core (A or B) internal/external memory and Peripheral Group 1; or between external memory and internal memory. The DCB1 bus runs at the system clock frequency (SCLK domain).

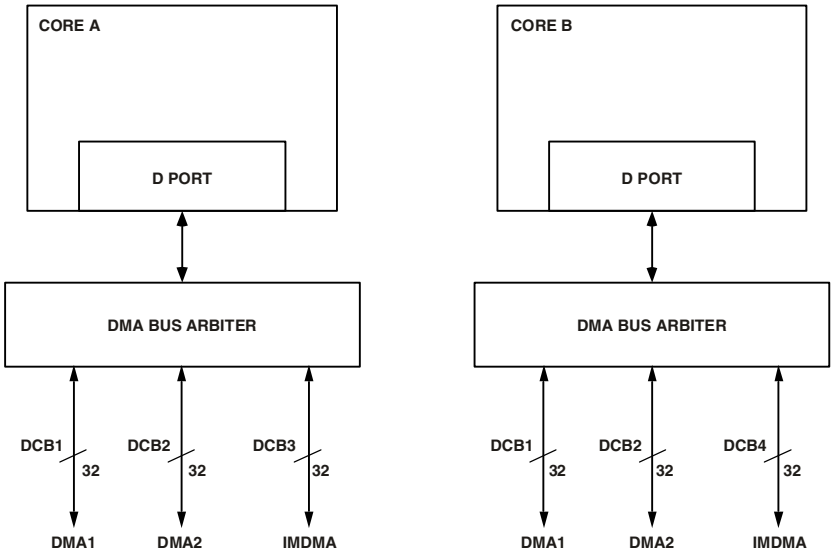

Figure 7-4. Core (A or B) Interface to DMA Controllers

- The DCB2 bus supports DMA data transfer between Core (A or B) internal/external memory and Peripheral Group 2; or between external memory and internal memory. The DCB2 bus runs at the system clock frequency (SCLK domain).
- The DCB3 and DCB4 buses support IMDMA data transfer between DSP Core (A or B) L1 memory and L2 memory. The DCB3 and DCB4 buses run at the core clock frequency (CCLK domain).

Since D Port access requests can come from three independent DMA controllers, DMA bus arbitration is necessary to resolve possible D Port access conflicts from DMA1, DMA2 and IMDMA. The D Port interface performs DMA bus (DCB1, DCB2 and DCB3) arbitration, converts transactions on these buses to the Core DMA bus protocol and conducts transactions via the Core DMA buses to L1 memory. Table 7-2 on page 7-12 describes the fixed priority arbitration scheme for the DMA

core access buses. Its configuration is defined by the CADMAPRIO/CBDMAPRIO bit of the SICA_SYSCR register.

Table 7-1. Core DMA Port DCB Bus Arbitration[1]

DMA controllers	Priority (CxDMAPRIO = 1)	Priority (CxDMAPRIO = 0)
DMA1	1	2
DMA2	2	1
IMDMA	3	3

1 In CxDMAPRIO, x = A for Core A and x = B for Core B.

On-Chip L2 Memory Interface

The L2 SRAM Memory block is organized into eight banks that can be accessed by either of two independent buses: the 64-bit Core L2 bus or the 32-bit Sys L2 bus. Figure 7-5 shows this interface diagram. L2 memory is organized as a multi-bank architecture of single-ported SRAMs, such that multiple accesses can occur in parallel, as long as they are to different banks. L2 memory has two ports: the Core L2 Port is connected to the 64-bit Core L2 bus and dedicated to Core (A or B) access requests and IMDMA access requests. The Sys L2 Port is connected to the 32-bit Sys L2 bus and dedicated to system DMA access requests. Two different banks can be accessed simultaneously by the 64-bit Core L2 bus and the 32-bit Sys L2 bus. When both buses attempt to access the same bank at the same time, the L2 memory arbitration logic resolves the conflict.

An L2 memory access requires two CCLK cycles for the access itself, plus any latency involved in the operation (see Table 7-3 on page 7-13). L2 memory interface control logic is clocked at the core frequency (CCLK clock domain). The system DMA access request comes from the DCB1 bus and the DCB2 bus, which run at system clock frequency (SCLK domain). The interface circuit synchronizes the DCB buses to the core clock domain and converts them to Sys L2 bus protocol. The IMDMA

access request comes from the DCB3 bus and the DCB4 bus, which run at the core clock frequency (CCLK domain).

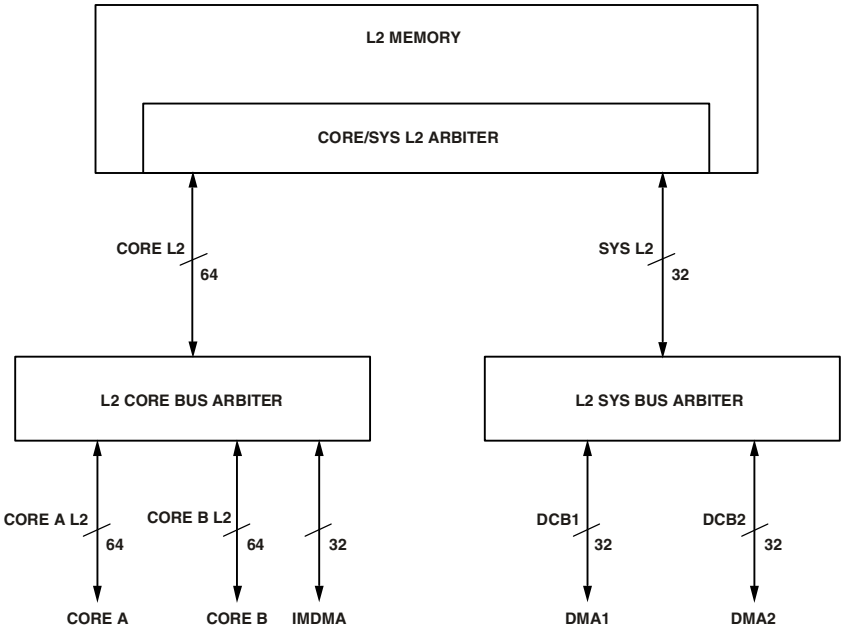

Figure 7-5. L2 Memory Bus Interfaces

As shown in Figure 7-5, there are several arbitration stages in the interface:

- Arbitration for Core L2 port access requests is based on a fixed priority scheme. Core A, Core B, and IMDMA, in order from highest to lowest priority, are the granted requestors of the Core L2 port. There is no arbitration while the current Core L2 bus requestor is performing locked or Cache Line Fill transactions.

  Once either Core A, Core B or IMDMA is granted the Core L2 bus, no other user can access this bus until the data transaction is accepted by L2 memory. For example, if IMDMA has requested an

L2 transaction and is waiting due to a bank conflict, a high priority Core A request that arrives subsequently will not be serviced until IMDMA wins the bank arbitration and sends its request to memory.

- Arbitration for Sys L2 port access request is based on a programmable priority scheme. After reset, the DCB2 bus (DMA2) has priority over the DCB1 bus (DMA1) for L2 memory accesses through the Sys L2 bus. However, this priority can be changed by programming the `L2DMAPRIO` bit in the `SICA_SYSCR` register.
- When both the Core L2 bus and the Sys L2 bus attempt to access the same bank at the same time, bank arbitration is required. Table 7-2 shows the L2 memory access bus arbitration.

Table 7-2. L2 Memory Interface Bus Arbitration

Requestor	Priority (L2DMAPRIO = 1)	Priority (L2DMAPRIO = 0)
Currently locked core access	1	1
Complete current core cache access	2	2
DMA1	3	4
DMA2	4	3
Core A L2	5	5
Core B L2	6	6
IMDMA	7	7

Table 7-3 shows the target latency and throughput for various types of accesses. Since the DMA bus has a dedicated port to the L1 and L2 memories, as long as the Core (A or B) access and DMA access are not to the

same memory bank, no stalls occur. DMA access to L1 or L2 memory can only be stalled by:

- An access already in progress from another DMA controller/channel.
- A core access already in progress, which locks the bank to be addressed.

For more details about DMA data transfer latency and DMA traffic control/optimization, please see "DMA Performance Optimization" on page 9-43.

Table 7-3. L2 Memory Interface Data Transaction Latency

Transaction type	Number of cycles to complete
Core L2 Read	9 CCLKs for each read
Dual DAG Read (same instruction)	9 CCLKs (first 32-bit fetch) 2 CCLKs (second 32-bit fetch)
Cache Line Fill (data and instruction)	9 CCLKs (first 64-bit fetch) 2-2-2 CCLKs (for next three 64-bit fetches)
Dual DAG Cache Line Miss (same instruction)	9-2-2-2 CCLKs (first miss, four 64-bit fetches) 2-2-2-2 CCLKs (second miss, four 64-bit fetches)
64-bit Instruction Fetch	9 CCLKs
IMDMA Read	2 CCLKs
Sys DMA Read	1 SCLK plus 2 CCLKs
Sys DMA Write	1 SCLK

When executing code from L2 memory, a core can fetch a 64-bit word. In the best case, the 64-bit word would contain four 16-bit instructions. For consecutive fetches of single cycle, 16-bit instructions, the maximum execution rate is four instructions every nine CCLKs—due to pre-fetching by the core. If core A and core B attempt to fetch an instruction in the same cycle, core A always has priority.

When the core (A or B) writes to L2 memory, a write buffer within the interface of each core improves performance. Up to five writes can be made to L2 memory without stalling a core. The sixth write, and subsequent writes when the buffer is full, take four CCLKs for each write. Specifically, a loop of eight writes to L2 memory would take five CCLKs for the first five writes plus four CCLKs for each of the three subsequent writes.

Peripheral Access Bus (PAB)

The ADSP-BF561 has a dedicated peripheral access bus (PAB) that connects all off-core peripherals to System MMR registers. The low-latency peripheral access bus keeps core stalls to a minimum and allows for manageable interrupt latencies to time-critical peripherals. All peripheral resources accessed through the PAB bus are mapped into the System MMR space of the ADSP-BF561 memory map. Each core can access System MMR space through the PAB bus. Although each core processor has byte addressability, the programming model is restricted to only 32-bit (aligned) access to the System MMRs. Byte access to this region is not supported.

PAB Arbitration

Figure 7-6 shows PAB bus connection to Cores A and B. System MMR access requests from Core A or B occur on the core DA0 bus. Each core has its own MMR bus controller that completes the clock domain and protocol conversions necessary to create a PAB access. The PAB arbiter controls which of the two MMR ports has access to the PAB. A simple arbitration policy is supported. After reset, Core A is granted the PAB bus. The currently granted requestor always has the highest priority. For example, if Core B is granted the PAB bus, Core A can not access it until the Core B data transaction is finished. One SCLK cycle is needed to change which core controls the PAB bus.

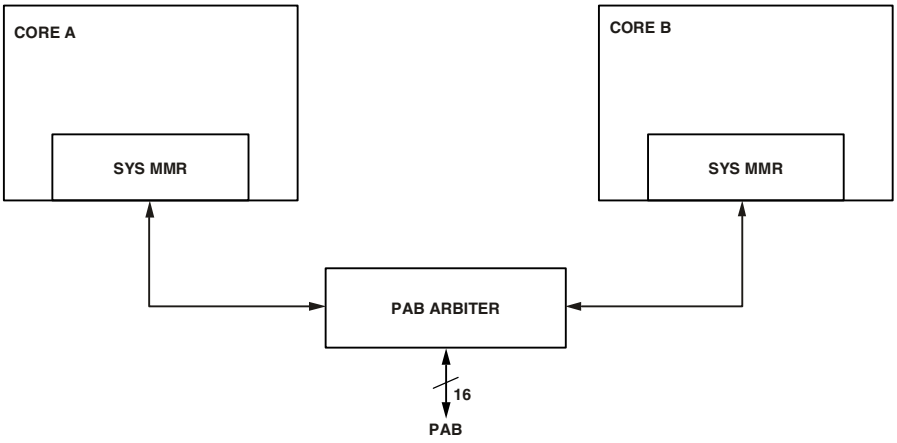

Figure 7-6. PAB Bus Arbitration Block Diagram

PAB Performance

For the PAB, the primary performance issue is latency, not throughput. Transfer latency for either read or write transfers on the PAB is two SCLK cycles. For example, the core can transfer up to 16 bits per access to the peripheral controllers connected with the PAB bus. With the core clock running at two times the frequency of the system clock, the first and subsequent System MMR read or write accesses take four core clock (CCLK) cycles of latency. With the core running three times the frequency of the system clock, the first and subsequent system MMR read or write accesses will take six core clock cycles of latency.

The following peripherals and function blocks have an interface to the PAB bus, which allows each processor core to access their control and status registers.

- Event Controller
- Clock and Power Management Controller
- Watchdog Timer

- General-Purpose Timers
- SPORT0
- SPORT1
- SPI
- Programmable Flags
- UART
- PPI0
- PPI1
- Asynchronous Memory Controller (AMC)
- SDRAM Controller (SDC)
- DMA Controllers

DMA-Related Buses

There are three types of DMA buses connected to three independent DMA controllers:

- DMA Access Buses (32-bit DAB1 and 16-bit DAB2) provide DMA data transfer to and from peripherals.
- DMA Core Buses (32-bit DCB1, DCB2, DCB3 and DCB4) provide DMA data transfer to and from Core (A or B) L1 memory, or DMA data transfer between L1 and L2 memory.
- DMA External Bus (DEB) provides DMA data transfer to and from external memory.

Figure 7-7 shows the DMA bus connections. These buses run at the system clock frequency (SCLK domain), except for the DCB3 and DCB4 buses, which run at the core clock frequency (CCLK domain).

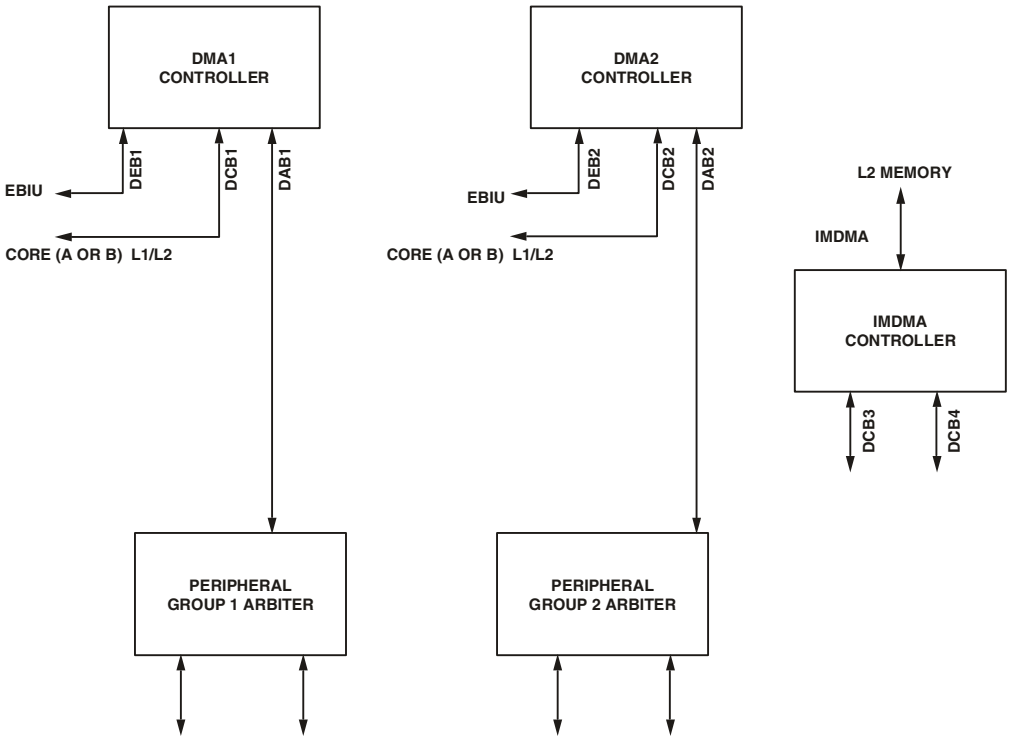

Figure 7-7. DMA Bus Connection and Arbitration Block Diagram

The 32/16-bit DAB bus provides DMA between the peripherals and L1/L2 internal memory through the DCB bus, or between peripherals and external memories through the DEB bus. A central DMA controller keeps track of DMA addresses and mediates the transfers. DMA is handled identically for 8-, 16- and 32-bit data sizes. The maximum bandwidth for any individual 16-bit peripheral is one 16-bit word transferred for every two

SCLK cycles. Peripherals that are capable of 32-bit DMA (and also configured for 32-bit mode) will transfer one 32-bit word every two SCLK cycles.

Peripheral DMA

There are nine DMA-capable peripherals in the ADSP-BF561 system managed by DMA controller 1 (DMA1) and DMA controller 2 (DMA2). Each DMA controller also has four memory DMA channels for DMA data transfer between external memory and L1 or L2 memory. There is another dedicated internal memory DMA controller (IMDMA) that supports DMA data transfer between L1 and L2 memory. The peripheral DMA controllers can transfer data between peripherals and internal or external memory.

The DCB bus arbitration for L2 configured as SRAM is shown in Table 7-2 on page 7-12. For off-chip memory, the access priority is shown in Table 7-5 on page 7-22. The ADSP-BF561 has programmable priority for peripherals on the DAB bus. For details about programmable DMA peripheral and DMA channel mapping, see "Direct Memory Access" on page 9-1.

DAB, DCB, and DEB Performance

The ADSP-BF561 DAB buses support 8-bit, 16-bit and 32-bit data transfers. DAB1 is a 32-bit data bus. DAB2 is a 16-bit bus. Both operate at the system clock rate, at a maximum frequency of 133 MHz, although a single peripheral DMA channel on a DAB bus operates at a maximum of SCLK/2. The DCB buses have a dedicated D Port into L1 memory and another dedicated Sys L2 port into L2 memory. No stalls occur as long as the core access and the DMA access are not to the same memory bank. If there is a conflict, DMA is the highest priority requester, followed by the core. Note that a locked transfer by the core processor (for example, execution of a TESTSET instruction) effectively disables arbitration for the addressed memory bank or resource until the memory lock is deasserted.

DMA controllers cannot perform locked transfers. DMA access to L1 memory can only be stalled by an access already in progress from another DMA channel.

Memory DMA transfers can result in repeated accesses to the same memory location. Because the memory DMA controller has the potential of simultaneously accessing on-chip and off-chip memory, considerable throughput can be achieved. The throughput rate for an on-chip/off-chip memory access is limited by the slower of the two accesses.

In the case where the transfer is from on-chip to on-chip memory or from off-chip to off-chip memory, the burst accesses cannot occur simultaneously. The transfer rate is then determined by adding each transfer plus an additional cycle between each transfer.

Table 7-4 shows many types of 32-bit memory DMA transfers. In the table, it is assumed that no other DMA activity is conflicting with ongoing operations. The numbers in the table are theoretical values. These values may be higher when they are measured on actual hardware due to a variety of reasons relating to the device that is connected to the EBIU.

For non-DMA accesses (for example, a core access via the EAB), a 32-bit access to SDRAM (of the form `R0 = [P0];` where P0 points to an address in SDRAM) will always be more efficient than executing two 16-bit accesses (of the form `R0 = W[P0++];` where P0 points to an address in SDRAM). In this example, a 32-bit SDRAM read will take 10 `SCLK` cycles while 2 16-bit reads will take 9 `SCLK` cycles each.

Table 7-4. Performance of DMA Access to External Memory

Source	Destination	Approximate SCLKs For n Words (from start of DMA to interrupt at end)
32-bit SDRAM	L1 Data memory	n + 14
L1 Data memory	32-bit SDRAM	n + 11

Table 7-4. Performance of DMA Access to External Memory (Cont'd)

Source	Destination	Approximate SCLKs For n Words (from start of DMA to interrupt at end)
32-bit Async memory	L1 Data memory	xn + 12, where x is the number of wait states + setup/hold SCLK cycles (minimum x = 2)
L1 Data memory	32-bit Async memory	xn + 9, where x is the number of wait states + setup/hold SCLK cycles (minimum x = 2)
32-bit SDRAM	32-bit SDRAM	10 + (17n/7)
32-bit Async memory	32-bit Async memory	10 + 2xn, where x is the number of wait states + setup/hold SCLK cycles (minimum x = 2)
L1 Data memory	L1 Data memory	2n + 12

The EAB buses and the DEB buses must arbitrate for access to external memory through the EBIU. Figure 7-8 shows the bus connection to the EBIU and the bus arbiters. Users must manage specific memory access traffic patterns to ensure that isochronous peripherals have enough allocated bandwidth and appropriate maximum data latency for both internal and external memory accesses.

External Access Bus (EAB)

The External Access Bus (EAB) provides a way for the DSP core and the Memory DMA controller to directly access off-chip memory and high-throughput memory-to-memory DMA transfers. The EAB supports single-word accesses of either 8-bit, 16-bit or 32-bit data types. The EAB operates at the system clock rate, up to 133 MHz for most system applications.

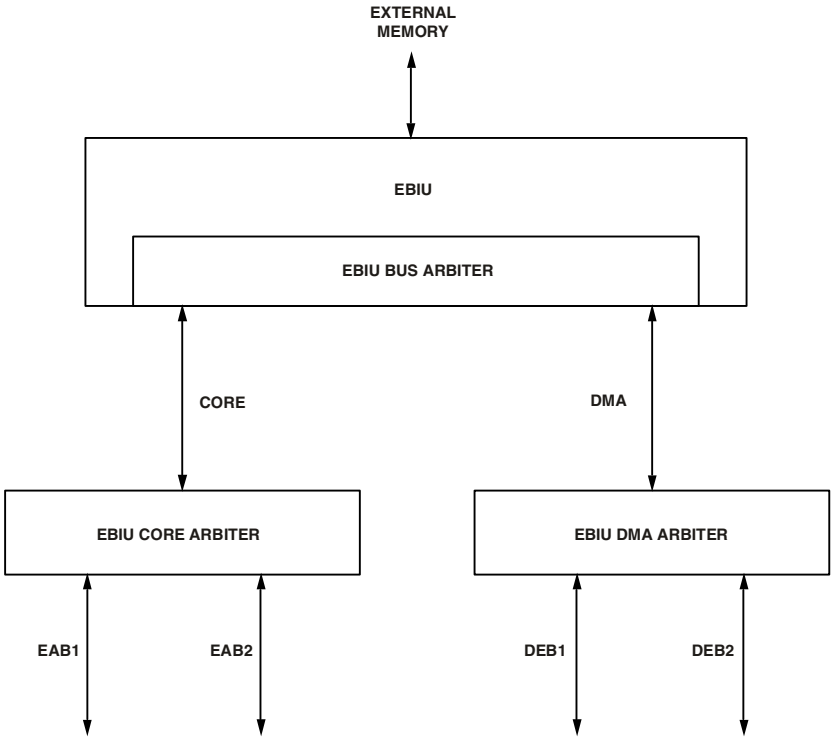

Figure 7-8. PAB Bus Arbitration Block Diagram

EAB Bus and DMA Bus Access Arbitration to EBIU

Use of bus interface resources to access the external memory port requires arbitration due to possible contention between the buses. As shown in Figure 7-8, there are three arbiters in the interface between Core A or B and the EBIU, and in the interface between the DMA controllers (DMA1 and DMA2) and the EBIU. A programmable priority arbitration scheme is used. By default, core accesses to the EBIU will be higher priority than non-urgent DMA requests. Core A EBIU accesses are higher priority than Core B accesses. When the `CDPRIO` bit of the `EBIU_AMGCTL` register is set, all DMA1 and DMA2 EBIU accesses are elevated to “urgent” status and will

have higher priority than Core A and Core B accesses. The external memory controller uses the EXTDMAPRIO bit of the SICA_SYSCR register to configure the relative priority of DMA1 and DMA2 with respect to each other, as shown in Table 7-5.

Table 7-5. EBIU Access Bus Arbitration

Requestor	Priority			
	(CDPRIO=0 and EXTDMAPRIO=1)	(CDPRIO=0 and EXTDMAPRIO=0)	(CDPRIO=1 and EXTDMAPRIO=1)	(CDPRIO=1 and EXTDMAPRIO=0)
Currently locked core access	1	1	1	1
Complete current core cache access or instruction fetch	2	2	2	2
Urgent DMA 1 (DEB1)	3	4	3	4
Urgent DMA 2 (DEB2)	4	3	4	3
Core A EAB1	5	5	5	5
Core B EAB2	6	6	6	6
DMA1 (DEB1)	7	8	When CDPRIO = 1, all DMA1 and DMA2 transfers are urgent.	
DMA2 (DEB2)	8	7		

For more information regarding the EAB bus bandwidth and data transaction latency, please see “External Bus Interface Unit” on page 16-1.

8 DYNAMIC POWER MANAGEMENT

This chapter describes the Dynamic Power Management functionality of the Blackfin processor. This functionality includes:

- Clocking
- Phase-Locked Loop (PLL)
- Dynamic Power Management Controller
- Operating Modes
- Voltage Control

Although this is a dual core processor, there is a single on-chip PLL module, and each core runs at the same core clock frequency (CCLK) and in the same power mode.

Clocking

The input clock into the processor, CLKIN, provides the necessary clock frequency, duty cycle, and stability to allow accurate internal clock multiplication by means of an on-chip Phase-Locked-Loop (PLL) module. During normal operation, the user programs the PLL with a multiplication factor for CLKIN. The resulting, multiplied signal is the Voltage-Controlled Oscillator (VCO) clock. A user-programmable value then divides the VCO clock signal to generate the core clock (CCLK).

A user-programmable value divides the VCO signal to generate the system clock (SCLK). The SCLK signal that clocks the Peripheral Access Bus (PAB),

DMA Access Bus (DAB) External Access Bus (EAB) and External Bus Interface Unit (EBIU).

These buses run at the PLL frequency divided by 1-15 (SCLK domain). Using the SSEL parameter of the PLL Divide register, select a divider value that allows these buses to run at or below the maximum SCLK rate specified in the *ADSP-BF561 Embedded Processor Data Sheet*.

To optimize performance and power dissipation, the processor allows the core and system clock frequencies to be changed dynamically in a "coarse adjustment." For a "fine adjustment," the PLL clock frequency can also be varied.

Phase-Locked Loop and Clock Control

To provide the clock generation for the core and system, the processor uses an analog PLL with programmable state-machine control.

The PLL design serves a wide range of applications. It emphasizes embedded and portable applications and low-cost, general purpose processors, in which performance, flexibility and control of power dissipation are key features. This broad range of applications requires a wide range of frequencies for the clock generation circuitry. The input clock may be a crystal, a crystal oscillator, or a buffered, shaped clock derived from an external system clock oscillator.

The PLL interacts with the Dynamic Power Management Controller (DPMC) block to provide power management functions for the processor. For information about the DPMC, see "Dynamic Power Management Controller" on page 8-11.

PLL Overview

Subject to the maximum VCO frequency, the PLL supports a wide range of multiplier ratios and achieves multiplication of the input clock, CLKIN.

To achieve this wide multiplication range, the processor uses a combination of programmable dividers in the PLL feedback circuit and output configuration blocks.

Figure 8-1 illustrates a conceptual model of the PLL circuitry, configuration inputs and resulting outputs. In the figure, the VCO is an intermediate clock from which the core clock (CCLK) and system clock (SCLK) are derived.

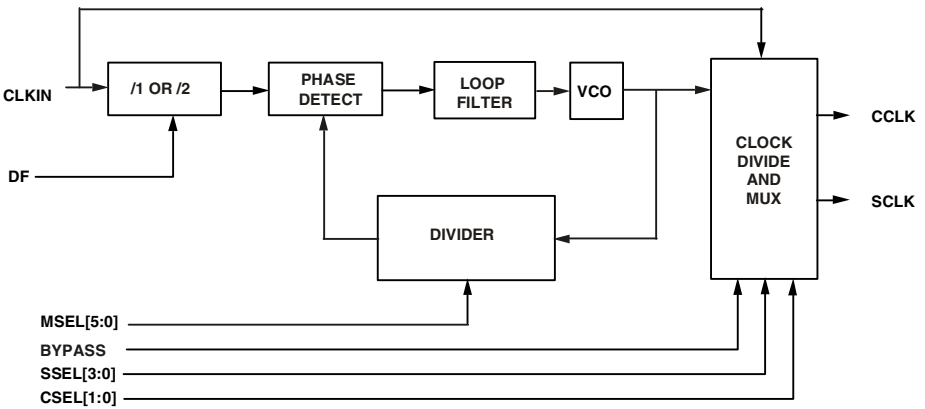

Figure 8-1. PLL Block Diagram

PLL Clock-Multiplier Ratios

The PLL Control register (PLL_CTL) governs the operation of the PLL. For details about the PLL_CTL register, see "PLL Control Register (PLL_CTL)" on page 8-7.

The Divide Frequency (DF) bit and Multiplier Select (MSEL[5:0]) field configure the various PLL clock dividers:

- DF enables the input divider.
- MSEL[5:0] controls the feedback dividers.

- The reset value of MSEL is 0xA. This value can be reprogrammed at startup in the boot code.

Table 8-1 illustrates the VCO multiplication factors for the various MSEL and DF settings.

As shown in the table, different combinations of MSEL[5:0] and DF can generate the same VCO frequencies. For a given application, one combination may provide lower power or satisfy the VCO maximum frequency. Under normal conditions, setting DF to 1 typically results in lower power dissipation. See the *ADSP-BF561 Embedded Processor Data Sheet* for maximum and minimum frequencies for CLKIN, CCLK, and VCO.

Table 8-1. MSEL Encodings

Signal Name MSEL[5:0]	VCO Frequency	
	DF=0	DF=1
0	64x	32x
1	1x	0.5x
2	2x	1x
N=3-62	Nx	0.5Nx
63	63x	31.5x

Core Clock/System Clock Ratio Control

Table 8-2 describes the programmable relationship between the VCO frequency and the core clock. Table 8-3 shows the relationship of the VCO frequency to the system clock. Note that the divider ratio must be chosen to limit SCLK to a frequency as specified in the *ADSP-BF561 Embedded Processor Data Sheet*. SCLK drives all synchronous, system-level logic.

The divider-ratio control bits, CSEL and SSEL, are in the PLL Divide register (PLL_DIV). For information about this register, see "PLL Divide Register (PLL_DIV)" on page 8-6.

The reset value of CSEL[1:0] is 0x0, and the reset value of SSEL[3:0] is 0x5. These values can be reprogrammed at startup by the boot code.

By writing the appropriate value to PLL_DIV, you can change the CSEL and SSEL value dynamically. Note that the divider ratio of the core clock can never be greater than the divider ratio of the system clock. If the PLL_DIV register is programmed to illegal values, the SCLK divider is automatically increased to be greater than or equal to the core clock divider.

The PLL_DIV register can be programmed at any time to change the CCLK and SCLK divide values without entering the Idle state. As long as the MSEL and DF control bits in the PLL Control register (PLL_CTL) remain constant, the PLL is locked.

Table 8-2. Core Clock Ratio

Signal Name CSEL[1:0]	Divider Ratio VCO/CCLK	Example Frequency Ratios (MHz)	
		VCO	CCLK
00	1	300	300
01	2	600	300
10	4	600	150
11	8	400	50

If changing the clock ratio by writing a new SSEL value into PLL_DIV, take care that the enabled peripherals do not suffer data loss due to SCLK frequency changes.

Table 8-3. System Clock Ratio

Signal Name SSEL[3:0]	Divider Ratio VCO/SCLK	Example Frequency Ratios (MHz)	
		VCO	SCLK
0000	Reserved	N/A	N/A
0001	1:1	133	133
0010	2:1	200	100

Table 8-3. System Clock Ratio (Cont'd)

Signal Name SSEL[3:0]	Divider Ratio VCO/SCLK	Example Frequency Ratios (MHz)	
		VCO	SCLK
0011	3:1	400	133
0100	4:1	500	125
0101	5:1	600	120
0110	6:1	600	100
N=7-15	N:1	600	600/N

PLL Memory-Mapped Registers (MMRs)

The user interface to the PLL is through four MMRs:

- The PLL Divide Register (PLL_DIV)
- The PLL Control register (PLL_CTL)
- The PLL Status register (PLL_STAT)
- The PLL Lock Count register (PLL_LOCKCNT)

All four registers are 16-bit MMRs and must be accessed with aligned 16-bit reads/writes.

PLL Divide Register (PLL_DIV)

The PLL Divide register (PLL_DIV) divides the PLL output clock to create the Core Clock (CCLK) and the System Clock (SCLK). (See Figure 8-2.) These values can be independently changed during processing to reduce power dissipation without changing the PLL state. The only restrictions are that the resulting CCLK frequency must be greater than or equal to the SCLK frequency, and SCLK must fall within the allowed range specified in the *ADSP-BF561 Embedded Processor Data Sheet*.

If the CCLK and SCLK divide values are programmed otherwise, the SCLK value is automatically adjusted to be slower than or equal to the core clock. Figure 8-2 shows the bits in the PLL_DIV register.

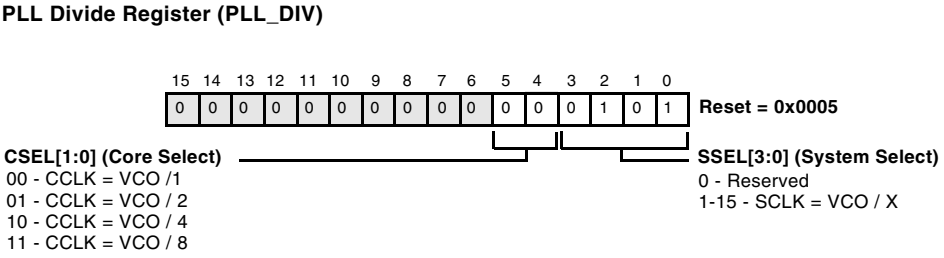

Figure 8-2. The PLL Divide Register

PLL Control Register (PLL_CTL)

The PLL Control register (PLL_CTL) controls operation of the PLL (see Figure 8-3). Note that changes to the PLL_CTL register do not take effect immediately. In general, the PLL_CTL register is first programmed with new values, and then a specific PLL programming sequence must be executed to implement the changes. See "PLL Programming Sequence" on page 8-19.

The following fields of the PLL_CTL register are used to control the PLL:

- SPORT_HYS – This bit is used to add 250 mV of hysteresis to the SPORT input pins to provide better immunity to system noise on SPORT clock and frame sync signals configured as inputs.
- MSEL[5:0] – The Multiplier Select (MSEL) field defines the input-clock-to-VCO-clock ((CLKIN-to-VCO) multiplier.
- BYPASS – The BYPASS bit is used to bypass the PLL. When BYPASS is set, CLKIN is passed directly to the core and peripheral clocks.

PLL Control Register (PLL_CTL)

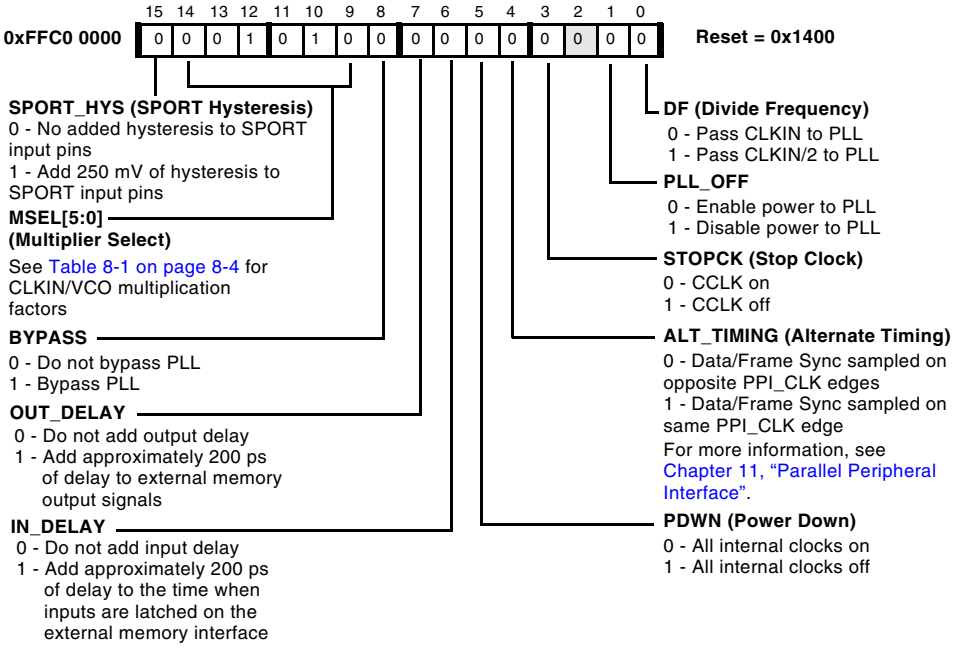

Figure 8-3. The PLL Control Register

- Output Delay – This bit is used to add approximately 200ps of delay to external memory output signals. See the *ADSP-BF561 Embedded Processor Data Sheet* to determine if this bit should be set.
- Input Delay – This bit is used to add approximately 200ps of delay to the time when inputs are latched on the external memory interface. See the *ADSP-BF561 Embedded Processor Data Sheet* to determine if this bit should be set.

- PDWN – The Power Down (PDWN) bit is used to place the processor in the Deep Sleep operating mode.
- For information about operating modes, see "Operating Modes" on page 8-12.
- STOPCK – The Stop Clock (STOPCK) bit is used to enable/disable the core clock, CCLK.
- PLL_OFF – This bit is used to enable/disable power to the PLL.
- DF – The Divide Frequency (DF) bit determines whether CLKIN is passed directly to the PLL or CLKIN/2 is passed.

PLL Status Register (PLL_STAT)

The PLL Status register (PLL_STAT) indicates the operating mode of the PLL and processor (see Figure 8-4). For more information about operating modes, see "Operating Modes" on page 8-12.

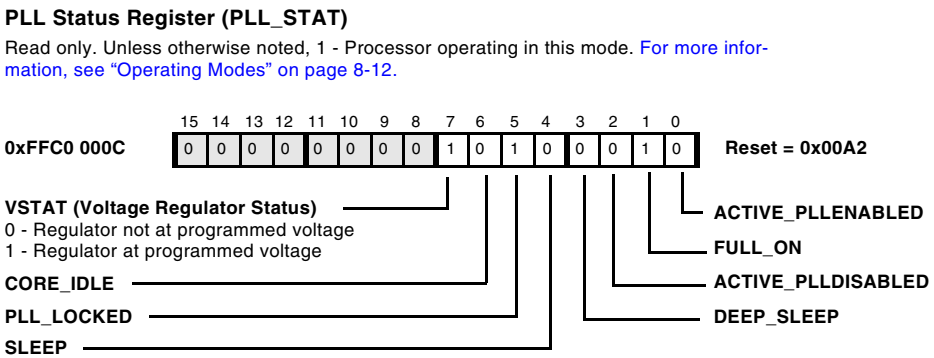

Figure 8-4. PLL Status Register

The following fields are used in the PLL_STAT register:

- VSTAT (Voltage Regulator Status) – This bit indicates whether the voltage regulator has reached the programmed voltage.
- When changing voltage levels, both cores must be put into an Idle operating state to allow the PLL to lock with the new voltage level. See "PLL Programming Sequence" on page 8-19.
- CORE_IDLE – This field is set to 1 when the Blackfin core is idled; that is, an IDLE instruction has executed, and the core is awaiting a wake-up signal.
- PLL_LOCKED – This field is set to 1 when the internal PLL lock counter has incremented to the value set in the PLL Lock Count register (PLL_LOCKCNT). For more information, see "PLL Lock Count Register (PLL_LOCKCNT)".
- SLEEP – This field is set to 1 when the processor is in the Sleep operating mode.
- DEEP_SLEEP – This field is set to 1 when the processor is in the Deep Sleep operating mode.
- ACTIVE_PLLDISABLED – This field is set to 1 when the processor is in the Active operating mode with the PLL powered down.
- FULL_ON – This field is set to 1 when the processor is in the Full On operating mode.
- ACTIVE_PLLENABLED – This field is set to 1 when the processor is in the Active operating mode with the PLL powered up.

PLL Lock Count Register (PLL_LOCKCNT)

When changing clock frequencies in the PLL, the PLL requires time to stabilize and lock to the new frequency.

The PLL Lock Count register (PLL_LOCKCNT) defines the number of CLKIN cycles that occur before the processor sets the PLL_LOCKED bit in the PLL_STAT register. When executing the PLL programming sequence, the internal PLL lock counter begins incrementing upon execution of the IDLE instruction. The lock counter increments by 1 each CLKIN cycle. When the lock counter has incremented to the value defined in the PLL_LOCKCNT register, the PLL_LOCKED bit is set.

See the *ADSP-BF561 Embedded Processor Data Sheet* for more information about PLL stabilization time and values that should be programmed into this register. For more information about operating modes, see "Operating Modes" on page 8-12. For more about the PLL programming sequence, see "PLL Programming Sequence" on page 8-19.

PLL Lock Count Register (PLL_LOCKCNT)

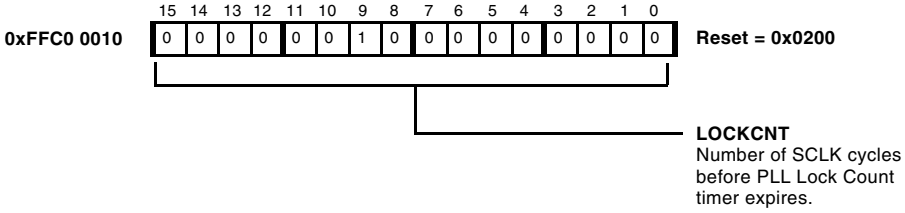

Figure 8-5. PLL Lock Count Register

Dynamic Power Management Controller

The Dynamic Power Management Controller (DPMC) of the processor works in conjunction with the PLL, allowing the user to control the processor's performance characteristics and power dissipation dynamically.

The DPMC provides the following features that allow the user to control performance and power:

- Multiple operating modes – The processor has four operating modes, each with different performance characteristics and power dissipation profiles. See "Operating Modes" on page 8-12.
- Peripheral clocks – Clocks to each peripheral are disabled automatically when the peripheral is disabled.
- Voltage control – The processor provides an on-chip switching regulator controller which, with some external components, can generate internal voltage levels from the external V_{DD} (V_{DDEXT}) supply.
- Depending on the needs of the system, the voltage level can be reduced to save power. See "Voltage Regulator Control Register (VR_CTL)" on page 8-24.

Operating Modes

The processor operates in four operating modes, each with unique performance and power-saving benefits. Table 8-4 summarizes the operational characteristics of each mode.

Table 8-4. Operational Characteristics

Operating Mode	Power Savings	PLL		CCLK	SCLK	Allowed DMA Access
		Status	Bypassed			
Full On	None	Enabled	No	Enabled	Enabled	L1/L2
Active	Medium	Enabled [1]	Yes	Enabled	Enabled	L1/L2
Sleep	High	Enabled	No	Disabled	Enabled	–
Deep Sleep	Maximum	Disabled	–	Disabled	Disabled	–

1 PLL can also be disabled in this mode.

Dynamic Power Management Controller States

Power management states are synonymous with the PLL control state. The state of the DPMC/PLL can be determined by reading the PLL Status register (see "PLL Status Register (PLL_STAT)" on page 8-9). In all modes except Sleep and Deep Sleep, the core can either execute instructions or be in Idle core state. If the core is in the Idle state, it can be awakened.

In all modes except Active, the SCLK frequency is determined by the SSEL-specified ratio to VCO. In Sleep mode, although the core clock is disabled, SCLK continues to run at the specified SSEL ratio.

The following sections describe the DPMC/PLL states in more detail, as they relate to the power management controller functions.

The processor starts operation in Full On or Active mode after the hardware reset is released, dependent on the state of the Bypass control pin (distinguish from Bypass bit). This pin is latched during reset and is used only to control reset related transitions.

Bypass pin = 0 Full On mode

Bypass pin = 1 Active mode

The Bypass pin has no effect on transitions from Sleep or between Active and Full On modes; which are controlled by the Bypass bit of the PLL_CTL register. See "Sleep Mode" on page 8-14, Figure 8-3 on page 8-8, and Figure 8-6 on page 8-16.

Full On Mode

Full On is the maximum performance mode of the processor. In this mode, the PLL is enabled and not bypassed. Full On is the normal execution state of the processor, with the processor and all enabled peripherals running at full speed. DMA access is available to L1 memories. From the

Full On mode, the processor can transition directly to the Active, Sleep or Deep Sleep modes, as shown in Figure 8-6 on page 8-16.

Active Mode

In the Active mode, the PLL is enabled but bypassed. Because the PLL is bypassed, the processor's core clock (CCLK) and system clock (SCLK) run at the input clock (CLKIN) frequency. DMA access is available to appropriately configured L1 and L2 memories.

In the Active mode, it is possible not only to bypass, but also to disable the PLL. If disabled, the PLL must be re-enabled before transitioning to the Full On or Sleep modes.

From the Active mode, the processor can transition directly to the Full On, Sleep, or Deep Sleep modes.

Sleep Mode

The Sleep mode significantly reduces power dissipation by idling the processor core. CCLK is disabled in this mode; however, SCLK continues to run at the speed configured by the MSEL and SSEL bit settings. Since CCLK is disabled, DMA access is available only to external memory in the Sleep mode. From the Sleep mode, a wake-up event causes the processor to transition to one of the following modes:

- Active mode if the BYPASS bit in the PLL_CTL register is set
- Full On mode if the BYPASS bit is cleared

Deep Sleep Mode

The Deep Sleep mode maximizes power-savings by disabling the PLL, CCLK and SCLK. In this mode, the processor core and all peripherals are disabled. DMA is not supported in this mode.

The Deep Sleep mode can be exited only by a hardware-reset event. A hardware reset begins the hardware reset sequence. For more information about hardware reset, see "Hardware Reset" on page 3-13.

When in the Deep Sleep operating mode, clocking to the SDRAM is turned off. Before entering the Deep Sleep mode, software should ensure that SDRAM has entered in self-refresh mode or that important information in SDRAM has been saved to a non-volatile memory.

When in either Sleep or Deep Sleep mode, the SLEEP pin will change to its active state. For more details, please refer to the *ADSP-BF561 Embedded Processor Data Sheet*.

Hibernate State

For lowest possible power dissipation, this state allows the internal supply (V_{DDINT}) to be powered down, while keeping the I/O supply (V_{DDEXT}) running. Although not strictly an operating mode like the four modes detailed above, it is illustrative to view it as such in the diagram of Figure 8-6. Since this feature is coupled to the on-chip switching regulator controller, it is discussed in detail in "Powering Down the Core (Hibernate State)" on page 8-28.

Operating Mode Transitions

Figure 8-6 graphically illustrates the operating modes and allowed transitions among them. In the diagram, ellipses represent operating modes. Arrows between the ellipses show the allowed transitions into and out of each mode.

The text next to each transition arrow shows the fields in the PLL Control Register (PLL_CTL) that must be changed for the transition to occur. For example, the transition from the Full On mode to Sleep mode indicates that the STOPCK bit must be set to 1 and the PDWN bit must be set to 0.

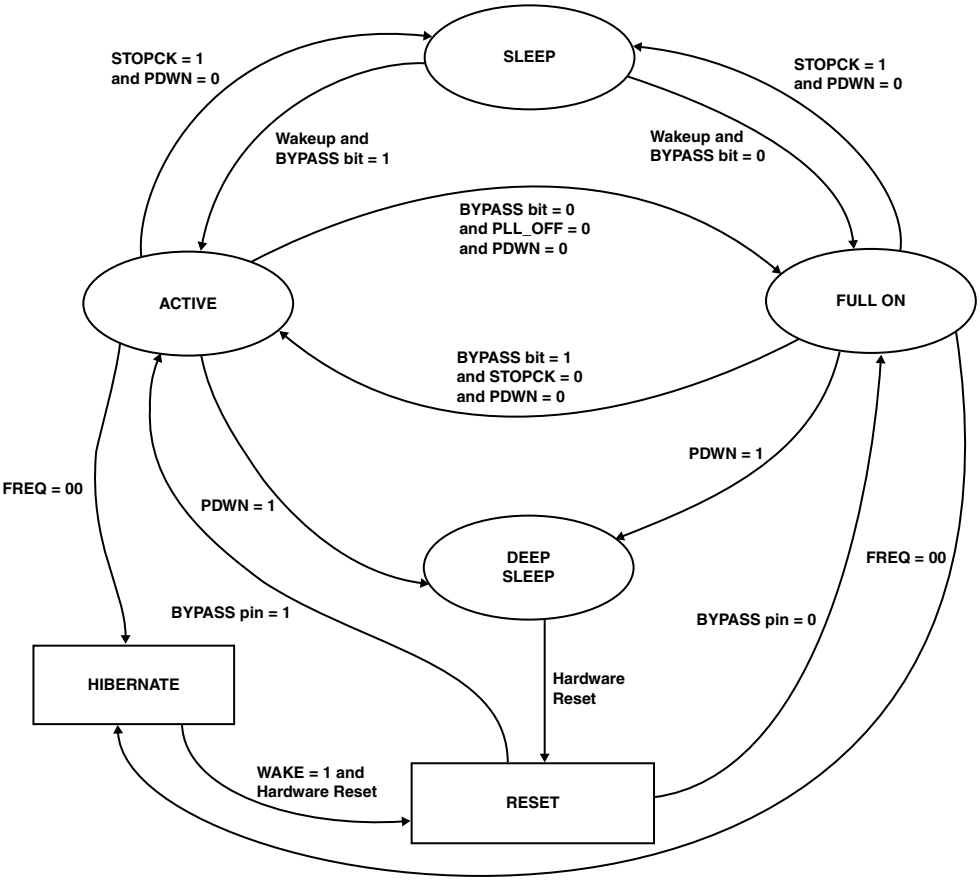

Figure 8-6. Operating Mode Transitions

In addition to the mode transitions shown in Figure 8-6, the PLL can be modified while in the Active operating mode. Power to the PLL can be applied and removed, and new clock-in-to-VCO-clock (CLKIN-to-VCO) multiplier ratios can be programmed.

Described in detail below, these changes to the PLL do not take effect immediately. As with operating mode transitions, the PLL-programming sequence must be executed for these changes to take effect (see "PLL Programming Sequence" on page 8-19).

- PLL Disabled: In addition to being bypassed in the Active mode, power to the PLL can be removed.
- When power is removed from the PLL, additional power savings are achieved, although they are relatively small. To remove power to the PLL, set the PLL_OFF bit in the PLL_CTL register, and then execute the PLL programming sequence.
- PLL Enabled: When the PLL is powered down, power can be reapplied later when additional performance is required.
- Power to the PLL must be reapplied before transitioning to the Full On or Sleep operating modes. To apply power to the PLL, clear the PLL_OFF bit in the PLL_CTL register, and then execute the PLL programming sequence.
- New Multiplier Ratio in Active Mode: New clock-in-to-VCO-clock (CLKIN-to-VCO) multiplier ratios can be programmed while in the Active mode.
- Although the CLKIN-to-VCO multiplier changes are not realized in the Active mode, forcing the PLL to lock to the new ratio in the Active mode before transitioning to Full On reduces the transition time, because the PLL is already locked to the new ratio. Note that the PLL must be powered up to lock to the new ratio. To program a new CLKIN-to-VCO multiplier, write the new MSEL[5:0] and/or DF values to the PLL_CTL register; then execute the PLL programming sequence.

- New Multiplier Ratio in Full On Mode: The multiplier ratio can also be changed while in Full On mode.
- In this case, the PLL state automatically transitions to the Active mode while the PLL is locking. After locking, the PLL returns to the Full On mode. To program a new CLKIN-to-VCO multiplier, write the new MSEL[5:0] and/or DF values to the PLL_CTL register; then execute the PLL programming sequence (on page 8-19).

Table 8-5 summarizes the allowed operating mode transitions.

Attempting to cause mode transitions other than those shown in Table 8-5 causes unpredictable behavior.

Table 8-5. Allowed Operating Mode Transitions

New Mode	Current Mode			
	Full On	Active	Sleep	Deep Sleep
Full On	–	Allowed	Allowed	–
Active	Allowed	–	Allowed	Allowed
Sleep	Allowed	Allowed	–	–
Deep Sleep	Allowed	Allowed	–	–

Programming Operating Mode Transitions

The operating mode of the processor is defined by the state of the PLL_OFF, BYPASS, STOPCK and PDWN bits of the PLL Control register (PLL_CTL). Merely modifying the bits of the PLL_CTL register does not change the operating mode or the behavior of the PLL. Changes to the PLL_CTL register are realized only after executing a specific code sequence, which is shown in Listing 8-1. This code sequence first brings the processor to a known, idled state. Once in this idled state, the PLL recognizes and implements the changes made to the PLL_CTL register. After the changes take effect, the processor operates with the new settings, including the new operating mode, if one is programmed.

PLL Programming Sequence

If new values are assigned to MSEL or DF in the PLL Control register (PLL_CTL), the instruction sequence shown in Listing 8-1 puts those changes into effect. The PLL programming sequence is also executed when transitioning between operating modes.

In order to program the PLL, both cores must be in an idled state. Alternatively, it is acceptable for core B to be in the "SRAM init" state that it enters after processor reset. For more information about this state, refer to "Booting Methods" on page 3-19.

Changes to the divider-ratio bits, CSEL and SSEL, can be made dynamically; they do not require execution of the PLL programming sequence.

Listing 8-1. PLL Programming Sequence

```
CLI R0; /* disable interrupts */
IDLE; /* drain pipeline, and send core into IDLE state */
STI R0; /* re-enable interrupts after wakeup */
```

The first two instructions in the sequence take the core to an idled state with interrupts disabled; the interrupt mask (IMASK) is saved to the R0 register, and the instruction pipeline is halted. The PLL state machine then loads the PLL_CTL register changes into the PLL.

If the PLL_CTL register changes include a new CLKIN-to-VCO multiplier or the changes re-apply power to the PLL, the PLL needs to re-lock. To re-lock, the PLL lock counter is first cleared, and then it begins incrementing, once per SCLK cycle. After the PLL lock counter reaches the value programmed into the PLL Lock Count register (PLL_LOCKCNT), the PLL sets the PLL_LOCKED bit in the PLL Status register (PLL_STAT), and the PLL asserts the PLL Wakeup interrupt.

Depending on how the PLL_CTL register is programmed, the processor proceeds in one of the following four ways:

- If the PLL_CTL register is programmed to enter either the Active or Full On operating mode, the PLL generates a wakeup signal, and then the processor continues with the STI instruction in the sequence, as described in "PLL Programming Sequence Continues" on page 8-21.
- When the state change enters Full On mode from Active mode or Active from Full On, the PLL itself generates a wake-up signal that can be used to exit the idled core state. The wake-up signal is generated by the PLL itself or another peripheral, watchdog or other timer, or other source. For more information about events that cause the processor to wake-up from being idled, see "System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26.
- If the PLL_CTL register is programmed to enter the Sleep operating mode, the processor immediately transitions to the Sleep mode and waits for a wake-up signal before continuing.

  When the wake-up signal has been asserted, the instruction sequence continues with the STI instruction, as described in the section, "PLL Programming Sequence Continues" on page 8-21, causing the processor to transition to:

  – Active mode if BYPASS in the PLL_CTL register is set

  – Full On mode if the BYPASS bit is cleared
- If the PLL_CTL register is programmed to enter the Deep Sleep operating mode, the processor immediately transitions to the Deep Sleep mode and waits for a hardware reset signal:

  – A hardware reset causes the processor to execute the reset sequence, as described in "Hardware Reset" on page 3-13.

- If no operating mode transition is programmed, the PLL generates a wakeup signal, and the processor continues with the STI instruction in the sequence, as described in the following section.

PLL Programming Sequence Continues

The instruction sequence shown in Listing 8-1 then continues with the STI instruction. Interrupts are re-enabled, IMASK is restored and normal program flow resumes.

To prevent spurious activity, DMA should be suspended while executing this instruction sequence.

Examples

The following code examples illustrate how to effect various operating mode transitions. Some setup code has been removed for clarity, and the following assumptions are made:

- P0 points to the PLL Control register (PLL_CTL). P1 points to the PLL Divide register.
- The PLL wake-up interrupt is enabled as a wake-up signal.
- MSEL[5:0] and DF in PLL_CTL are set to (b#011111) and (b#0) respectively, signifying a CLKIN to VCO multiplier of 31x.

Active Mode to Full On Mode

Listing 8-2 provides code for transitioning from Active operating mode to Full On.

Listing 8-2. Transitioning from Active Mode to Full On Mode

```
CLI R2; /* disable interrupts, copy IMASK to R2 */
R1.L=0x3E00; /* clear BYPASS bit */
W[P0]=R1; /* and write to PLL_CTL */
```

```
IDLE; /* drain pipeline, enter idled state, wait for PLL wake-up
*/
STI R2; /* after PLL wakeup occurs, restore interrupts and IMASK
*/
… /* processor is now in Full On mode */
```

Full On Mode to Active Mode

Listing 8-3 provides code for transitioning from Full On operating mode to Active.

Listing 8-3. Transitioning from Full On Mode to Active Mode

```
CLI R2; /* disable interrupts, copy IMASK to R2 */
R1.L=0x3F00; /* set BYPASS bit */
W[P0]=R1; /* and write to PLL_CTL */

IDLE; /* drain pipeline, enter idled state, wait for PLL wake-up
*/
STI R2; /* after PLL wakeup occurs, restore interrupts and IMASK
*/
... /* processor is now in Active mode */
```

In the Full On Mode, Change CLKIN to VCO Multiplier From 31x to 2x

Listing 8-4 provides code for changing `CLKIN`-to-VCO multiplier from 31x to 2x in the Full On operating mode.

Listing 8-4. Changing the CLKIN to VCO Multiplier

```
CLI R2; /* disable interrupts, copy IMASK to R2 */
R1.L=0x0400; /* change VCO multiplier to 2x */
W[P0]=R1; /* by writing to PLL_CTL */
```

```
IDLE; /* drain pipeline, enter idled state, wait for PLL wake-up
*/
STI R2; /* after PLL wakeup occurs, restore interrupts and IMASK
*/
... /* processor is now in Full On mode, with the CLKIN to CCLK
multiplier set to 2x */
```

Dynamic Supply Voltage Control

In addition to clock frequency control, the processor provides the capability to run the core processor at different voltage levels. As power dissipation is proportional to the voltage squared, significant power reductions can be accomplished when lower voltages are used.

The processor uses two power domains. These power domains are shown in Table 8-6. Each power domain has a separate V_{DD} supply. Note that the internal logic of the processor and much of the processor I/O can be run over a range of voltages. See the *ADSP-BF561 Embedded Processor Data Sheet* for details on the allowed voltage ranges for each power domain and power-dissipation data.

Table 8-6. ADSP-BF561 Power Domains

Power Domain	V_{DD} Range
All internal logic	Variable
All other I/O	Variable

Power Supply Management

The processor provides an on-chip switching regulator controller which, with some external hardware, can generate internal voltage levels from the external V_{DDEXT} supply with an external power transistor as shown in

Figure 8-7. This voltage level can be reduced to save power, depending upon the needs of the system.

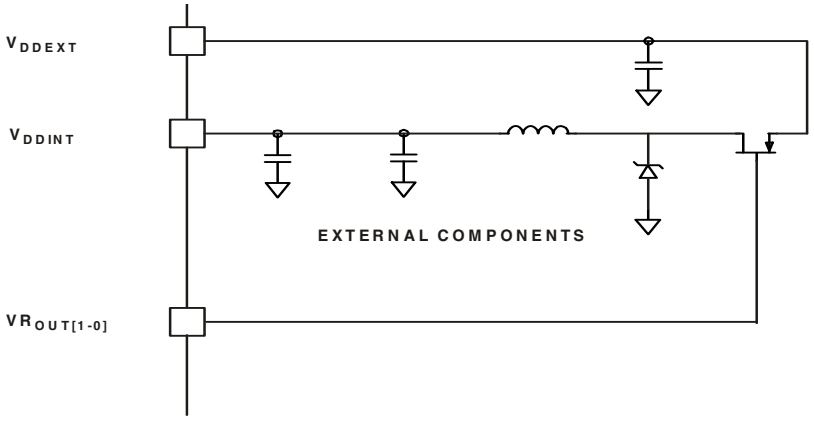

Figure 8-7. ADSP-BF561 Voltage Regulator

Voltage Regulator Control Register (VR_CTL)

The on-chip core voltage regulator controller manages the internal logic voltage levels for the V_{DDINT} supply. The Voltage Regulator Control register (VR_CTL) controls the regulator (see Figure 8-8).

Voltage Regulator Control Register (VR_CTL)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
0	0	0	0	0	0	0	0	1	1	0	1	1	0	1	1	**Reset = 0x00DB**

SCKELOW (Drive SCKE Low During Reset)
0 - Allow SCKE to go high (=1) during reset
1 - Maintain SCKE low (=0) during reset

WAKE (Wake-up Enable)
0 - Reset wake-up disabled.
1 - Reset wake-up enabled.

VLEV[3:0] (Internal Voltage Level)
See Table 8-7 for encodings.

FREQ[1:0] (Switching Frequency)
Controls the switching oscillator frequency for the voltage regulator. See Table 8-9 for encodings.

GAIN[1:0] (Voltage Level Gain)
Controls how quickly the voltage output settles on its final value. See Table 8-8 for encodings.

Figure 8-8. Voltage Regulator Control Register

Writing to VR_CTL initiates a PLL relock sequence, thus the PLL reprogramming sequence must be followed after modifying VR_CTL. For more information, see “PLL Programming Sequence” on page 8-19.

The following fields of the VR_CTL register are used to control internal logic voltage levels:

- SCKELOW – The Drive SCKE Low During Reset (SCKELOW) control bit protects against the default reset state behavior of setting the EBIU pins to their inactive state. This bit should be set if the SDRAM has been properly configured and is being placed into self-refresh mode while the processor is in Hibernate state. Failure to set this bit results in the SCKE pin going high during reset, which takes the SDRAM out of self-refresh mode, resulting in data decay in the SDRAM due to loss of refresh rate.
- WAKE – The Wake-up Enable (WAKE) control bit allows the voltage regulator to be awakened from powerdown (FREQ=00) upon an interrupt from a low-going edge on the $\overline{\text{RESET}}$ pin.
- VLEV[3:0] – The Voltage Level (VLEV) field identifies the nominal internal voltage level. Please refer to the *ADSP-BF561 Embedded Processor Data Sheet* for the applicable VLEV voltage range and associated voltage tolerances.
- FREQ[1:0] – The Frequency (FREQ) field controls the switching oscillator frequency for the voltage regulator. A higher frequency setting allows for smaller switching capacitor and inductor values, while potentially generating more EMI (electromagnetic interference).

To bypass onboard regulation, program a value of b#00 in the FREQ field, and leave the VROUT pins floating.

- GAIN[1:0] – The Gain (GAIN) field controls the internal loop gain of the switching regulator loop; this bit controls how quickly the voltage output settles on its final value. In general, higher gain allows for quicker settling times but causes more overshoot in the process.

Table 8-7 lists the voltage level values for VLEV[3:0].

Table 8-7. VLEV Encodings

VLEV	Voltage
0000-0101	Reserved
0110	0.85 volts
0111	0.90 volts
1000	0.95 volts
1001	1.00 volts
1010	1.05 volts
1011	1.10 volts
1100	1.15 volts
1101	1.20 volts
1110	1.25 volts
1111	1.30 volts

For legal VLEV values with respect to the voltage tolerance, consult the appropriate processor-specific data sheet.

Table 8-8 lists the switching frequency values for FREQ[1:0].

Table 8-8. FREQ Encodings

FREQ	Value
00	Powerdown/bypass onboard regulation
01	333KHz
10	667KHz
11	1MHz

Table 8-9 lists the gain levels configured by GAIN[1:0].

Table 8-9. Gain Encodings

Gain	Value
00	5
01	10
10	20
11	50

Changing Voltage

Minor changes in operating voltage can be accommodated without requiring special consideration or action by the application program. See the *ADSP-BF561 Embedded Processor Data Sheet* for more information about voltage tolerances and allowed rates of change.

Reducing the operating voltage to greatly conserve power or raising the operating voltage to greatly increase performance will probably require significant changes to the operating voltage level. To ensure predictable behavior when varying the operating voltage, the processor should be brought to a known and stable state before the operating voltage is modified.

The recommended procedure is to follow the PLL programming sequence when varying the voltage. After changing the voltage level in the VR_CTL register, the PLL will automatically enter the Active mode when both processors enter the Idle state. At that point the voltage level will change and the PLL will re-lock with the new voltage. After the PLL_LOCKCNT has expired, the part will return to the Full On state. When changing voltages, a larger PLL_LOCKCNT value may be necessary than when changing just the PLL frequency. See the *ADSP-BF561 Embedded Processor Data Sheet* for details.

After the voltage has been changed to the new level, the processor can safely return to any operational mode so long as the operating parameters, such as core clock frequency (CCLK), are within the limits specified in the *ADSP-BF561 Embedded Processor Data Sheet* for the new operating voltage level.

Powering Down the Core (Hibernate State)

The internal supply regulator for the processor can be shut off by writing b#00 to the FREQ bits of the VR_CTL register. This disables both CCLK and SCLK. Furthermore, it sets the internal power supply voltage (V_{DDINT}) to 0 V, eliminating any leakage currents from the processor. The internal supply regulator can be woken up by asserting the $\overline{\text{RESET}}$ pin.

If the on-chip supply controller is bypassed, so that V_{DDINT} is sourced externally, the only way to power down the core is to remove the external V_{DDINT} voltage source.

When the core is powered down, V_{DDINT} is set to 0 V, and thus the internal state of the processor is not maintained. Therefore, any critical information stored internally (memory contents, register contents, and so on) must be written to a non-volatile storage device prior to removing power.

In order to maintain the self-refresh mode of SDRAM during hibernate, the SCKE pin must be programmed. Read the discussion on page 8-25 for details on using the SCKELOW bit to program the SCKE pin.

Powering down V_{DDINT} does not affect V_{DDEXT}. While V_{DDEXT} is still applied to the processor, external pins are maintained at a three-state level, unless otherwise specified.

To power down the internal supply:

1. Write to VR_CTL, setting the FREQ bits to b#00, and the WAKE bit to 1.
2. Execute this code sequence:

```
CLI R0 ;

IDLE ;
```

3. When the Idle state is reached, V_{DDINT} will transition to 0 V.

Failure to allow V_{DDINT} to complete the transition to 0 V before waking up the processor can cause undesired results.

4. When the processor is woken up, the PLL relocks and the boot sequence defined by the BMODE[1:0] pin settings takes effect.

9 DIRECT MEMORY ACCESS

The processor uses Direct Memory Access (DMA) to transfer data within memory spaces or between a memory space and a peripheral. The processor can specify data transfer operations and return to normal processing while the fully integrated DMA controller carries out the data transfers independent of processor activity.

ADSP-BF561 has three independent DMA controllers, DMA1, DMA2 and IMDMA. The centralized DMA control configuration with variable sized data structure descriptors are capable of data transfers within memories or between memory and a peripheral. Each of DMA1 and DMA2 has twelve peripheral DMA channels and four memory DMA channels. IMDMA has four memory DMA channels.

Peripherals connected to the DMA1 controller can support up to 32-bit data transfers, while those connected to DMA2 can support up to 16-bit transfers. For more information, refer to Chapter 7, “Chip Bus Hierarchy”.

The DMA controllers can perform several types of data transfers:

- Between internal memories (L1/L2)(IMDMA, MDMA1, MDMA2) Chapter 6, “Memory”.
- Between external memory (SDRAM, Flash memory) and internal memory (L1/L2) (MDMA1, MDMA2) Chapter 6, “Memory”.
- Between memory and the Serial Peripheral Interface (SPI) Chapter 10, “SPI Compatible Port Controllers”.

- Between memory and a Serial Port (SPORT) Chapter 12, “Serial Port Controllers”.
- Between memory and the UART Port Chapter 13, “UART Port Controller”.
- Between memory and the Parallel Peripheral Interface (PPI) Chapter 11, “Parallel Peripheral Interface”.

Table 9-1 shows the default channel assignment at DSP reset. The channel number is also the default channel priority number, a value of 0 indicates the highest priority. However, the user can program it using the Peripheral Map Register (`DMAx_PERIPHERAL_MAP` or `MDMA_yy_PERIPHERAL_MAP`).

Table 9-1. Default Channel Assignment

DMA1	Description	DMA2	Description	IMDMA	Description
Channel0	PPI0 Receive/ Transmit	Channel0	SPORT0 Receive	Channel0	IMDMA Stream0 Destination
Channel1	PPI1 Receive/ Transmit	Channel1	SPORT0 Transmit	Channel1	IMDMA Stream0 Source
Channel2	Reserved	Channel2	SPORT1 Receive	Channel2	IMDMA Stream1 Destination
Channel3	Reserved	Channel3	SPORT1 Transmit	Channel3	IMDMA Stream1 Source
Channel4	Reserved	Channel4	SPI Receive/Transmit		
Channel5	Reserved	Channel5	UART Receive		
Channel6	Reserved	Channel6	UART Transmit		
Channel7	Reserved	Channel7	Reserved		
Channel8	Reserved	Channel8	Reserved		

Table 9-1. Default Channel Assignment (Cont'd)

DMA1	Description	DMA2	Description	IMDMA	Description
Channel9	Reserved	Channel9	Reserved		
Channel10	Reserved	Channel10	Reserved		
Channel11	Reserved	Channel11	Reserved		
Channel12	MDMA Stream0 Destination	Channel12	MDMA Stream0 Destination		
Channel13	MDMA Stream0 Source	Channel13	MDMA Stream0 Source		
Channel14	MDMA Stream1 Destination	Channel14	MDMA Stream1 Destination		
Channel15	MDMA Stream1 Source	Channel15	MDMA Stream1 Source		

This chapter describes the features common to all the DMA channels, as well as how DMA operations are set up. For specific peripheral features, see the appropriate peripheral chapter for additional information. Performance and bus arbitration for DMA operations can be found in "DAB, DCB, and DEB Performance" on page 7-18.

DMA transfers on the processor can be descriptor-based or register-based. Descriptor-based DMA transfers require a set of parameters stored within memory to initiate a DMA sequence. This sort of transfer allows the chaining together of multiple DMA sequences. In descriptor-based DMA operations, a DMA channel can be programmed to automatically set up and start another DMA transfer after the current sequence completes. Register-based DMA allows the processor to directly program DMA control registers to initiate a DMA transfer. On completion, the control registers may be automatically updated with their original setup values for continuous transfer, if needed.

The DMA engine in the ADSP-BF561 is a modified version of the DMA engine in the ADSP-BF533, which in turn has been completely revised

from that of the ADSP-BF535 in order to increase its flexibility and provide additional functionality.

DMA and Memory DMA MMRs

For convenience, discussions in this chapter use generic (non-peripheral specific) DMA and Memory DMA register names.

- Generic DMA register names are listed in Table 9-2.
- Generic Memory DMA register names are listed in Table 9-4 on page 9-9.

DMA registers fall into three categories:

- Parameter registers, such as `DMAx_y_CONFIG` and `DMAx_y_X_COUNT`

The letter ‘x’ in DMAx_y represents which DMA_1 or 2 is being referred. The letter ‘y’ represents a specific DMA-capable peripheral. For example, for DMA with default channel mapping, `DMA2_5_CONFIG` represents the `DMA_CONFIG` register for the UART *RX* (connected to channel 5 of DMA2) peripheral. For default DMA2 channel mappings, see Table 9-5 on page 9-25.

- Only Parameter registers can be loaded directly from descriptor elements; descriptor elements are listed in Table 9-3 on page 9-7.
- Current registers, such as `DMAx_y_CURR_ADDR` and `DMAx_y_CURR_X_COUNT`
- Control/Status registers, such as `DMAx_y_IRQ_STATUS` and `DMAx_y_PERIPHERAL_MAP`

Table 9-2 lists the generic names of the DMA registers. For each register, the table also shows the MMR offset, a brief description of the register, the register category, and reset value.

Table 9-2. Generic DMA Memory-Mapped Register Names

MMR Offset	Generic MMR Name	MMR Description	Register Category	Reset Value
0x00	NEXT_DESC_PTR	Link pointer to next descriptor	Parameter	0x0000 0000
0x04	START_ADDR	Start address of current buffer	Parameter	0x0000 0000
0x08	DMA_CONFIG	DMA Configuration register, including enable bit	Parameter	0x0000
0x0C	Reserved	Reserved		
0x10	X_COUNT	Inner loop count	Parameter	0x0001
0x14	X_MODIFY	Inner loop address increment, in bytes	Parameter	0x0002
0x18	Y_COUNT	Outer loop count (two-dimensional only)	Parameter	0x0001
0x1C	Y_MODIFY	Outer loop address increment, in bytes	Parameter	0x0002
0x20	CURR_DESC_PTR	Current Descriptor Pointer	Current	
0x24	CURR_ADDR	Current DMA Address	Current	
0x28	IRQ_STATUS	Interrupt Status register: Contains Completion and DMA Error Interrupt status and channel state (Run/Fetch/Paused)	Control/ Status	
0x2C	PERIPHERAL_MAP	Peripheral to DMA Channel Mapping: Contains a 4-bit value specifying the peripheral to associate with this DMA channel (Read-only for MDMA channels)	Control/ Status	
0x30	CURR_X_COUNT	Current count (one-dimensional) or intra-row X count (two-dimensional); counts down from X_COUNT	Current	
0x34	Reserved	Reserved		

Table 9-2. Generic DMA Memory-Mapped Register Names (Cont'd)

MMR Offset	Generic MMR Name	MMR Description	Register Category	Reset Value
0x38	CURR_Y_COUNT	Current row count (two-dimensional only); counts down from Y_COUNT	Current	
0x3C	Reserved	Reserved		

All DMA registers can be accessed as 16-bit entities. However, the following registers may also be accessed as 32-bit registers:

- NEXT_DESC_PTR
- START_ADDR
- CURR_DESC_PTR
- CURR_ADDR

When these four registers are accessed as 16-bit entities, only the lower 16 bits can be accessed.

Naming Conventions for DMA MMRs

Because confusion might arise between descriptor element names and generic DMA register names, this chapter uses the naming conventions in Table 9-3, where:

- The left column lists the generic name of the MMR, which is used when discussing the general operation of the DMA engine.

Note the generic names in the left column are not actually mapped to resources in the processor.

- The middle column lists the specific MMR name. Only specific MMR names are mapped to processor resources.

The letter ‘x’ in DMAx_y represents which DMA Controller (1 or 2) is being referenced. For example, for DMA with default channel mapping, `DMA2_3_IRQ_STATUS` represents the `DMA_IRQ_STATUS` register for the DMA2 channel 3.

- The channel number can be assigned by default or can be programmed. For the DMA channel numbers and the default peripheral mapping, see Table 9-5 on page 9-25.
- The last column lists the macro assigned to each descriptor element in memory.
- The macro name in the last column serves only to clarify the discussion of how the DMA engine operates.

Table 9-3. Naming Conventions: DMA MMRs and Descriptor Elements

Generic MMR Name	Specific MMR Name (x= DMA Controller Number, y = DMA Channel Number)	Name of Corresponding Descriptor Element in Memory
DMA_CONFIG	DMAx_y_CONFIG	DMACFG
NEXT_DESC_PTR	DMAx_y_NEXT_DESC_PTR	NDPH (upper 16 bits), NDPL (lower 16 bits)
START_ADDR	DMAx_y_START_ADDR	SAH (upper 16 bits), SAL (lower 16 bits)
X_COUNT	DMAx_y_X_COUNT	XCNT
Y_COUNT	DMAx_y_Y_COUNT	YCNT
X_MODIFY	DMAx_y_X_MODIFY	XMOD
Y_MODIFY	DMAx_y_Y_MODIFY	YMOD
CURR_DESC_PTR	DMAx_y_CURR_DESC_PTR	N/A
CURR_ADDR	DMAx_y_CURR_ADDR	N/A

Table 9-3. Naming Conventions: DMA MMRs and Descriptor Elements (Cont'd)

Generic MMR Name	Specific MMR Name (x= DMA Controller Number, y = DMA Channel Number)	Name of Corresponding Descriptor Element in Memory
CURR_X_COUNT	DMAx_y_CURR_X_COUNT	N/A
CURR_Y_COUNT	DMAx_y_CURR_Y_COUNT	N/A
IRQ_STATUS	DMAx_y_IRQ_STATUS	N/A
PERIPHERAL_MAP	DMAx_y_PERIPHERAL_MAP	N/A

Naming Conventions for Memory DMA Registers

The names of Memory DMA registers differ somewhat from the names of other DMA registers. Memory DMA streams cannot be reassigned to different channels, whereas the peripherals associated with DMA can be mapped to any DMA channel between 0 and 11.

Table 9-4 shows the naming conventions for Memory DMA registers. In each name, the letters *yy* have four possible values:

- S0, Memory DMA Source Stream 0
- D0, Memory DMA Destination Stream 0
- S1, Memory DMA Source Stream 1
- D1, Memory DMA Destination Stream 1

Table 9-4. Naming Conventions for Memory DMA Registers

Generic MMR Name	Memory DMA MMR Name (i = DMA Controller Number, yy = S0, S1, D0, or D1)	Name of Corresponding Descriptor Element in Memory
DMA_CONFIG	MDMAx_yy_CONFIG	DMACFG
NEXT_DESC_PTR	MDMAx_yy_NEXT_DESC_PTR	NDPH (upper 16 bits), NDPL (lower 16 bits)
START_ADDR	MDMAx_yy_START_ADDR	SAH (upper 16 bits), SAL (lower 16 bits)
X_COUNT	MDMAx_yy_X_COUNT	XCNT
Y_COUNT	MDMAx_yy_Y_COUNT	YCNT
X_MODIFY	MDMAx_yy_X_MODIFY	XMOD
Y_MODIFY	MDMAx_yy_Y_MODIFY	YMOD
CURR_DESC_PTR	MDMAx_yy_CURR_DESC_PTR	N/A
CURR_ADDR	MDMAx_yy_CURR_ADDR	N/A
CURR_X_COUNT	MDMAx_yy_CURR_X_COUNT	N/A
CURR_Y_COUNT	MDMAx_yy_CURR_Y_COUNT	N/A
IRQ_STATUS	MDMAx_yy_IRQ_STATUS	N/A
PERIPHERAL_MAP	MDMAx_yy_PERIPHERAL_MAP	N/A
DMA_CONFIG	IMDMA_yy_CONFIG	DMACFG
NEXT_DESC_PTR	IMDMA_yy_NEXT_DESC_PTR	NDPH (upper 16 bits), NDPL (lower 16 bits)
START_ADDR	IMDMA_yy_START_ADDR	SAH (upper 16 bits), SAL (lower 16 bits)
X_COUNT	IMDMA_yy_X_COUNT	XCNT
Y_COUNT	IMDMA_yy_Y_COUNT	YCNT
X_MODIFY	IMDMA_yy_X_MODIFY	XMOD
Y_MODIFY	IMDMA_yy_Y_MODIFY	YMOD
CURR_DESC_PTR	IMDMA_yy_CURR_DESC_PTR	N/A
CURR_ADDR	IMDMA_yy_CURR_ADDR	N/A

Table 9-4. Naming Conventions for Memory DMA Registers (Cont'd)

Generic MMR Name	Memory DMA MMR Name (i = DMA Controller Number, yy = S0, S1, D0, or D1)	Name of Corresponding Descriptor Element in Memory
CURR_X_COUNT	IMDMA_yy_CURR_X_COUNT	N/A
CURR_Y_COUNT	IMDMA_yy_CURR_Y_COUNT	N/A
IRQ_STATUS	IMDMA_yy_IRQ_STATUS	N/A
PERIPHERAL_MAP	IMDMA_yy_PERIPHERAL_MAP	N/A

Next Descriptor Pointer Register (DMAx_y_NEXT_DESC_PTR/ MDMA_yy_NEXT_DESC_PTR)

The NEXT_DESC_PTR register specifies where to look for the start of the next descriptor block when the DMA activity specified by the current descriptor block finishes (see Figure 9-1).

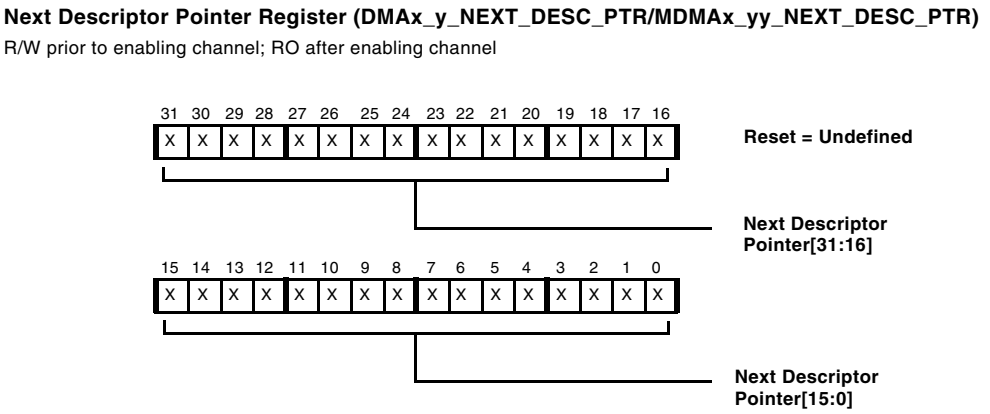

Figure 9-1. Next Descriptor Pointer Register

This register is used in small and large descriptor list modes. At the start of a descriptor fetch in either of these modes, the 32-bit NEXT_DESC_PTR register is copied into the CURR_DESC_PTR register. Then, during the descriptor fetch, the CURR_DESC_PTR register increments after each element of the descriptor is read in.

In small and large descriptor list modes, the NEXT_DESC_PTR register, and not the CURR_DESC_PTR register, must be programmed directly via MMR access before starting DMA operation.

In Descriptor Array mode, the Next Descriptor Pointer register is disregarded, and fetching is controlled only by the CURR_DESC_PTR register.

Start Address Register (DMAx_y_START_ADDR/ MDMAx_yy_START_ADDR)

The START_ADDR register, shown in Figure 9-2, contains the start address of the data buffer currently targeted for DMA.

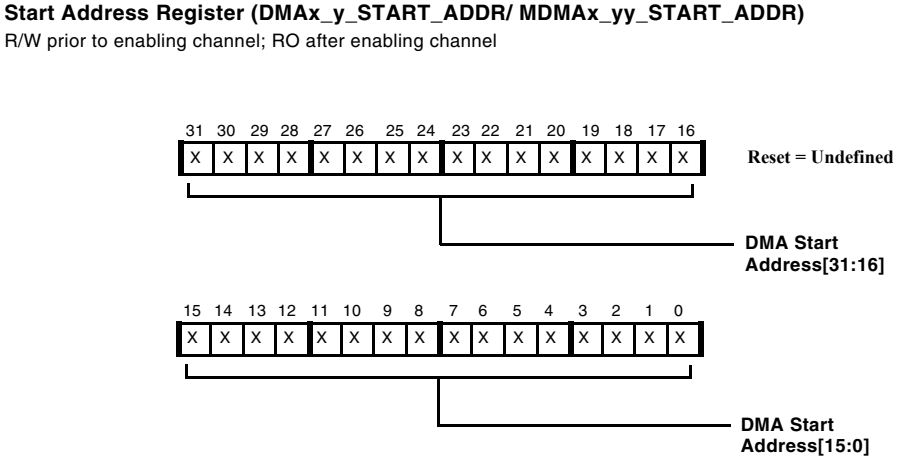

Figure 9-2. Start Address Register

DMA Configuration Register (DMAx_y_CONFIG/MDMAx_yy_CONFIG)

The DMA_CONFIG register, shown in Figure 9-3, is used to set up DMA parameters and operating modes. Note that writing the DMA_CONFIG register while DMA is already running will cause a DMA error unless writing with the DMAEN bit set to 0.

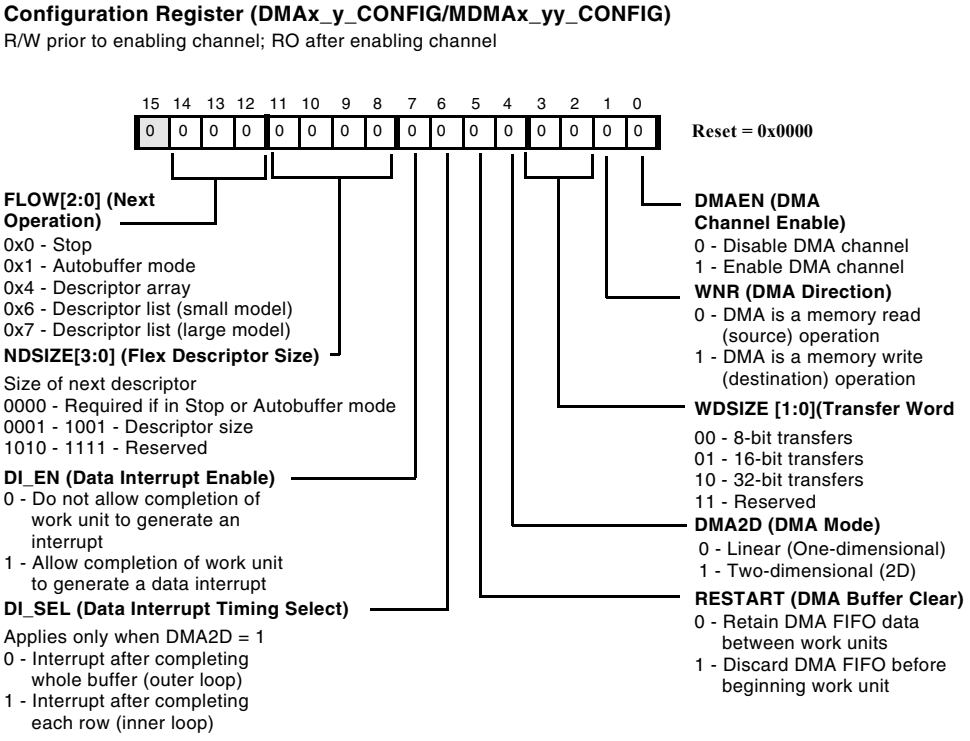

Figure 9-3. Configuration Register

The fields of the DMAx_y_CONFIG register are used to set up DMA parameters and operating modes.

- FLOW[2:0] (Next Operation). This field specifies the type of DMA transfer to follow the present one. The flow options are:
- 0x0 - Stop. When the current work unit completes, the DMA channel stops automatically, after signaling an interrupt (if selected). The DMA_RUN status bit in the DMAx_y_IRQ_STATUS register changes from 1 to 0, while the DMAEN bit in the DMAx_y_CONFIG register is unchanged. In this state, the channel is paused. Peripheral interrupts are still filtered out by the DMA unit. The channel may be restarted simply by another write to the DMAx_y_CONFIG register specifying the next work unit, where the DMAEN bit is set to 1.
- 0x1 - Autobuffer Mode. In this mode, no descriptors in memory are used. Instead, DMA is performed in a continuous circular buffer fashion based on user-programmed DMAx MMR settings. Upon completion of the work unit, the Parameter registers are reloaded into the Current registers, and DMA resumes immediately with zero overhead. Autobuffer mode is stopped by a user write of 0 to the DMAEN bit in the DMAx_y_CONFIG register.
- 0x4 - Descriptor Array Mode. This mode fetches a descriptor from memory that does not include the NDPH or NDPL elements. Because the descriptor does not contain a Next Descriptor Pointer entry, the DMA engine defaults to using the CURR_DESC_PTR register to step through descriptors, thus allowing a group of descriptors to follow one another in memory like an array.

- 0x6 - Descriptor List (Small Model) Mode. This mode fetches a descriptor from memory that includes NDPL, but not NDPH. Therefore, the high 16 bits of the Next Descriptor Pointer field are taken from the upper 16 bits of the NEXT_DESC_PTR register, thus confining all descriptors to a specific 64K page in memory.
- 0x7 - Descriptor List (Large Model) Mode. This mode fetches a descriptor from memory that includes NDPH and NDPL, thus allowing maximum flexibility in locating descriptors in memory.
- NDSIZE[3:0] (Flex Descriptor Size). This field specifies the number of descriptor elements in memory to load. This field must be 0 if in Stop or Autobuffer mode. If NDSIZE and FLOW specify a descriptor that extends beyond YMOD, a DMA error results.
- DI_EN (Data Interrupt Enable). This bit specifies whether to allow completion of a work unit to generate a data interrupt.
- DI_SEL (Data Interrupt Timing Select). This bit specifies the timing of a data interrupt—after completing the whole buffer or after completing each row of the inner loop. This bit is used only in 2D DMA operation.
- RESTART (DMA Buffer Clear). This bit specifies whether receive data held in the channel's data FIFO should be preserved (RESTART = 0) or discarded (RESTART = 1) before beginning the next work unit. Receive data is automatically discarded when the DMAEN bit changes from 0 to 1, typically when a channel is first enabled. Received FIFO data should usually be retained between work units if the work units make up a continuous datastream. If, however, a new work unit starts a new datastream, the RESTART bit should be set to 1 to clear out any previously received data.

The RESTART bit applies only to memory write DMA channels. It is reserved in the cases of memory read DMA channels and MDMA channels, and must be 0 in those cases.

In memory write DMA channels, the RESTART bit only affects the first work unit initiated by a write to the DMAx_y_CONFIG register. The RESTART bit has no effect if it is set in DMACFG elements of DMA descriptors.

- DMA2D (DMA Mode). This bit specifies whether DMA mode involves only X_COUNT and X_MODIFY (one-dimensional DMA) or also involves Y_COUNT and Y_MODIFY (two-dimensional DMA).
- WDSIZE[1:0] (Transfer Word Size). The DMA engine supports transfers of 8-, 16-, or 32-bit items. Each request/grant results in a single memory access. The DMA Address Pointer registers' increment sizes (strides) must be a multiple of the transfer unit size—1 for 8-bit, 2 for 16-bit, 4 for 32-bit.

Not all peripherals support al transfer sizes. For further information, refer to the chapter for each peripheral.

- WNR (DMA Direction). This bit specifies DMA direction—memory read (0) or memory write (1).
- DMAEN (DMA Channel Enable). This bit specifies whether to enable a given DMA channel.

When a peripheral DMA channel is enabled, interrupts from the peripheral denote DMA requests. When a channel is disabled, the DMA unit ignores the peripheral interrupt and passes it directly to the interrupt controller. To avoid unexpected results, take care to enable the DMA channel before enabling the peripheral, and to disable the peripheral before disabling the DMA channel.

Inner Loop Count Register (DMAx_y_X_COUNT/MDMAx_yy_X_COUNT)

For 2D DMA, the X_COUNT register, shown in Figure 9-4, contains the inner loop count. For 1D DMA, it specifies the number of elements to read in. For details, see "Two-Dimensional DMA" on page 9-38. A value of 0 in X_COUNT corresponds to 65,536 elements.

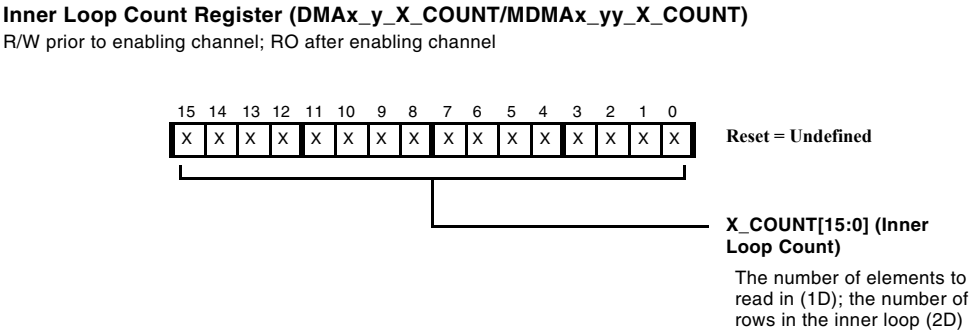

Figure 9-4. Inner Loop Count Register

Inner Loop Address Increment Register (DMAx_y_X_MODIFY/MDMAx_yy_X_MODIFY)

The Inner Loop Address Increment (X_MODIFY) register, shown in Figure 9-5, contains a signed, two's-complement byte-address increment. In 1D DMA, this increment is the stride that is applied after transferring each element.

Note X_MODIFY is specified in bytes, regardless of the DMA transfer size.

In 2D DMA, this increment is applied after transferring each element in the inner loop, up to but not including the last element in each inner loop. After the last element in each inner loop, the `Y_MODIFY` register is applied instead, except on the very last transfer of each work unit. The `X_MODIFY` register is always applied on the last transfer of a work unit.

The `X_MODIFY` field may be set to 0. In this case, DMA is performed repeatedly to or from the same address. This is useful, for example, in transferring data between a data register and an external memory-mapped peripheral.

Inner Loop Address Increment Register (DMAx_y_X_MODIFY/MDMAx_yy_X_MODIFY)
R/W prior to enabling channel; RO after enabling channel

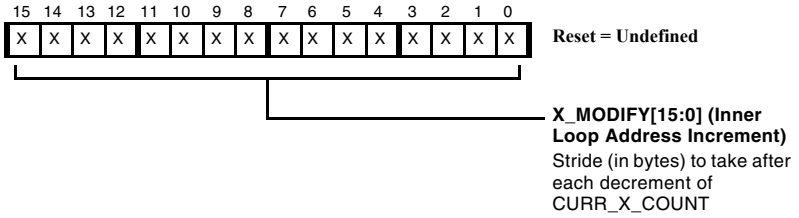

Figure 9-5. Inner Loop Address Increment Register

Outer Loop Count Register (DMAx_y_Y_COUNT/MDMAx_yy_Y_COUNT)

For 2D DMA, the Outer Loop Count (`Y_COUNT`) register, shown in Figure 9-6, contains the outer loop count. It is not used in 1D DMA mode. This register contains the number of rows in the outer loop of a 2D DMA sequence. For details, see "Two-Dimensional DMA" on page 9-38.

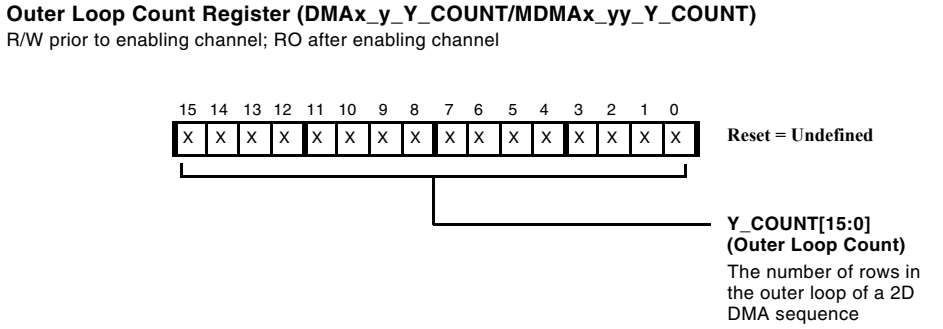

Figure 9-6. Outer Loop Count Register

Outer Loop Address Increment Register (DMAx_y_Y_MODIFY/MDMAx_yy_Y_MODIFY)

The Outer Loop Address Increment (`Y_MODIFY`) register (Figure 9-7) contains a signed, two's-complement value. This byte-address increment is applied after each decrement of the `CURR_Y_COUNT` register except for the last item in the 2D array where the `CURR_Y_COUNT` also expires. The value is the offset between the last word of one "row" and the first word of the next "row." For details, see "Two-Dimensional DMA" on page 9-38.

Note `Y_MODIFY` is specified in bytes, regardless of the DMA transfer size.

Outer Loop Address Increment Register (DMAx_y_Y_MODIFY/ MDMAx_yy_Y_MODIFY)
R/W prior to enabling channel; RO after enabling channel

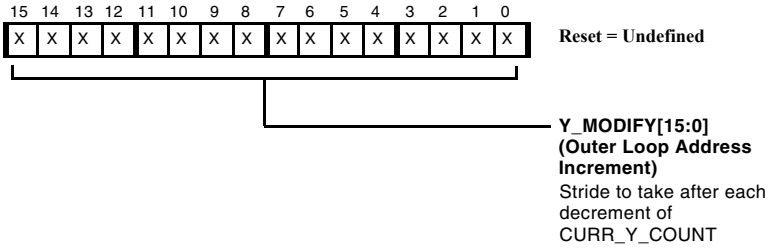

Figure 9-7. Outer Loop Address Increment Register

Current Descriptor Pointer Register (DMAx_y_CURR_DESC_PTR/ MDMAx_yy_CURR_DESC_PTR)

The Current Descriptor Pointer (CURR_DESC_PTR) register (Figure 9-8) contains the memory address for the next descriptor element to be loaded. For FLOW mode settings that involve descriptors (FLOW = 4, 6, or 7), this register is used to read descriptor elements into appropriate MMRs before a DMA work block begins. For Descriptor List modes (FLOW = 6 or 7), this register is initialized from the NEXT_DESC_PTR register before loading each descriptor. Then, the address in the CURR_DESC_PTR register increments as each descriptor element is read in.

When the entire descriptor has been read, the CURR_DESC_PTR register contains the value:

Descriptor Start Address + Descriptor Size (# of elements)

For Descriptor Array mode (FLOW = 4), this register (not the NEXT_DESC_PTR register) must be programmed by MMR access before starting DMA operation.

Current Descriptor Pointer Register (DMAx_y_CURR_DESC_PTR/ MDMAx_yy_CURR_DESC_PTR)

R/W prior to enabling channel; RO after enabling channel

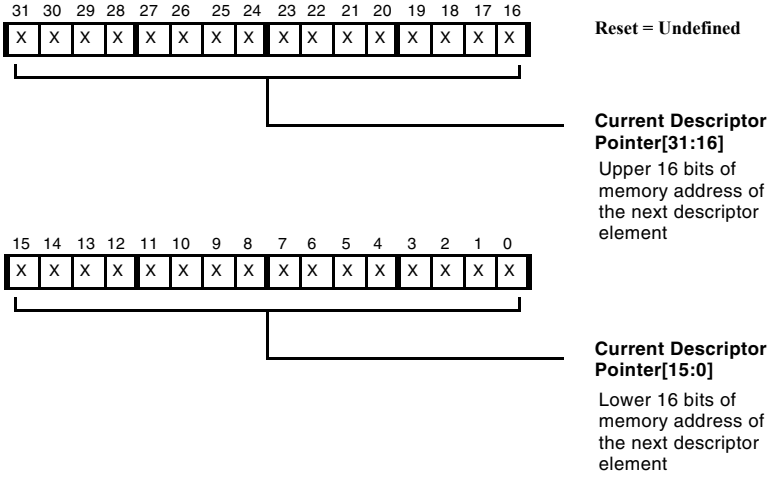

Figure 9-8. Current Descriptor Pointer Register

Current Address Register (DMAx_y_CURR_ADDR/ MDMAx_yy_CURR_ADDR)

The Current Address (CURR_ADDR) register, shown in Figure 9-9, contains the present DMA transfer address for a given DMA session. At the start of a DMA session, the CURR_ADDR register is loaded from the START_ADDR register, and it is incremented as each transfer occurs. The Current Address register contains 32 bits.

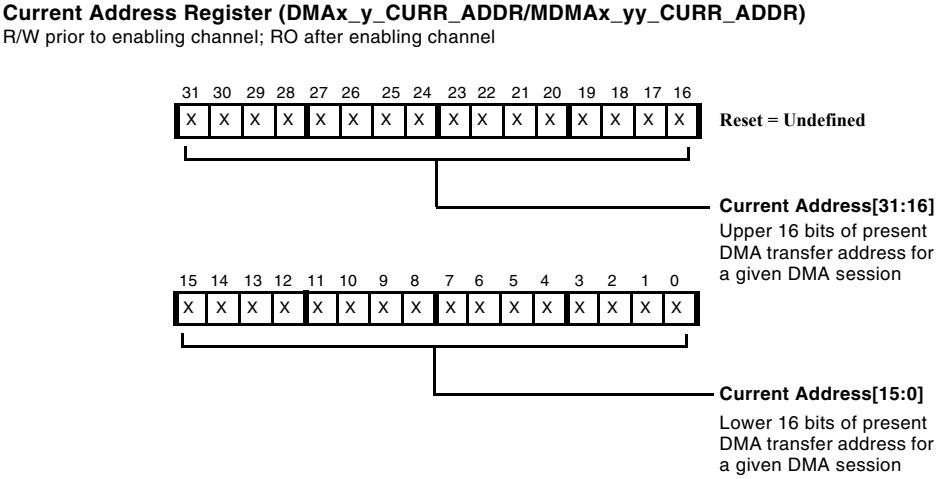

Figure 9-9. Current Address Register

Current Inner Loop Count Register (DMAx_y_CURR_X_COUNT/ MDMAx_yy_CURR_X_COUNT)

The Current Inner Loop Count (CURR_X_COUNT) register (Figure 9-10) is loaded by the X_COUNT register at the beginning of each DMA session (for 1D DMA) and also after the end of DMA for each row (for 2D DMA). Otherwise it is decremented each time an element is transferred. Expiration of the count in this register signifies that DMA is complete. In 2D DMA, the CURR_X_COUNT register value is 0 only when the entire transfer is complete. Between rows it is equal to the value of the X_COUNT register.

Current Inner Loop Count Register (DMAx_y_CURR_X_COUNT/ MDMAx_yy_CURR_X_COUNT)

R/W prior to enabling channel; RO after enabling channel

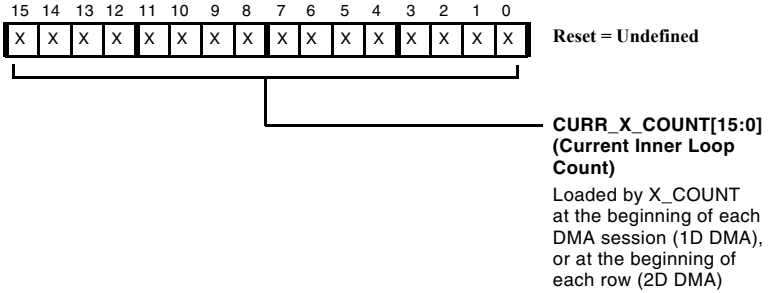

Figure 9-10. Current Inner Loop Count Register

Current Outer Loop Count Register (DMAx_y_CURR_Y_COUNT/ MDMAx_yy_CURR_Y_COUNT)

The Current Outer Loop Count (CURR_Y_COUNT) register (Figure 9-11) is loaded by the Y_COUNT register at the beginning of each 2D DMA session.

Current Outer Loop Count Register (DMAx_y_CURR_Y_COUNT/ MDMAx_yy_CURR_Y_COUNT)

R/W prior to enabling channel; RO after enabling channel

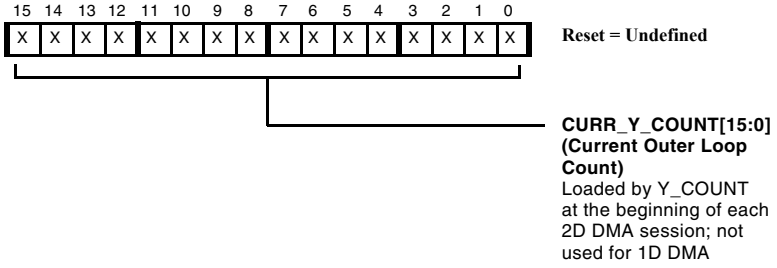

Figure 9-11. Current Outer Loop Count Register

It is not used for 1D DMA. This register is decremented each time the CURR_X_COUNT register expires during 2D DMA operation (1 to X_COUNT or 1 to 0 transition), signifying completion of an entire row transfer. After a 2D DMA session is complete, CURR_Y_COUNT = 1 and CURR_X_COUNT = 0.

Peripheral Map Register (DMAx_y_PERIPHERAL_MAP/ MDMAx_yy_PERIPHERAL_MAP)

Each DMA channel's PERIPHERAL_MAP register (Figure 9-12 and Figure 9-13) contains bits that:

- Map the channel to a specific peripheral.
- Identify whether the channel is a Peripheral DMA channel or a Memory DMA channel.

Note a 1:1 mapping should exist between DMA channels and peripherals. The user is responsible for ensuring that multiple DMA channels are not mapped to the same peripheral and that multiple peripherals are not mapped to the same DMA port. If multiple channels are mapped to the same peripheral, only one channel is connected (the lowest priority channel). If a nonexistent peripheral (for example, 0xF in the PMAP field) is mapped to a channel, that channel is disabled—DMA requests are ignored, and no DMA grants are issued. The DMA requests are also not forwarded from the peripheral to the interrupt controller.

Follow these steps to swap the DMA channel priorities of two channels. Assume that channels 6 and 7 are involved.

1. Make sure DMA is disabled on channels 6 and 7.
2. Write `DMA6_PERIPHERAL_MAP` with `0x7000` and `DMA7_PERIPHERAL_MAP` with `0x6000`.
3. Enable DMA on channels 6 and/or 7.

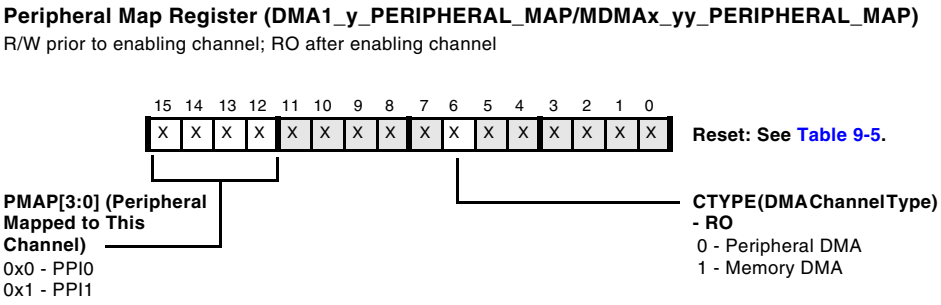

Figure 9-12. Peripheral Map Register

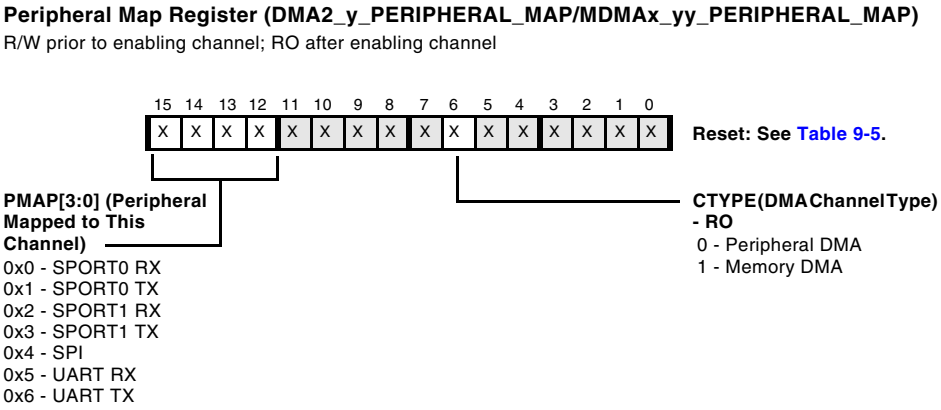

Figure 9-13. Peripheral Map Register

Table 9-5 lists the binary peripheral map settings for each DMA-capable peripheral.

Table 9-5. Peripheral Mapping

DMA Channel	DMA Unit	Default Peripheral Mapping	Default PERIPHERAL_MAP Setting	Comments
0 (highest priority)	2	SPORT0 RX	0x0000	
1	2	SPORT0 TX	0x1000	
2	2	SPORT1 RX	0x2000	
3	2	SPORT1 TX	0x3000	
4	2	SPI	0x4000	
5	2	UART RX	0x5000	
6	2	UART TX	0x6000	
7-11	2		0x7000-0xb000	
12-13	2	DMA2 Memory stream0	0x0040	Not re-assignable
14,15 (lowest priority)	2	DMA2 Memory stream1	0x0040	Not re-assignable
0 (highest priority)	1	PPI0	0x0000	
1	1	PPI1	0x1000	
2-11	1		0x2000-0xb000	
12,13	1	DMA1 Memory stream0	0x0040	Not re-assignable
14,15 (lowest priority)	1	DMA1 Memory stream1	0x0040	Not re-assignable

Interrupt Status Register (DMAx_y_IRQ_STATUS/MDMAx_yy_IRQ_STATUS)

The Interrupt Status (IRQ_STATUS) register, shown in Figure 9-14, contains bits that record whether the DMA channel:

- Is enabled and operating, enabled but stopped, or disabled.
- Is fetching data or a DMA descriptor.
- Has detected that a global DMA interrupt or a channel interrupt is being asserted.
- Has logged occurrence of a DMA error.

Note the DMA_DONE interrupt is asserted when the last memory access (read or write) has completed.

For a memory transfer to a peripheral, there may be up to four data words in the channel's DMA FIFO when the interrupt occurs. At this point, it is normal to immediately start the next work unit. If, however, the application needs to know when the final data item is actually transferred to the peripheral, the application can test or poll the DMA_RUN bit. As long as there is undelivered transmit data in the FIFO, the DMA_RUN bit is 1.

For a memory write DMA channel, the state of the DMA_RUN bit has no meaning after the last DMA_DONE event has been signaled. It does not indicate the status of the DMA FIFO.

For MemDMA transfers where it is not desired to use an interrupt to notify when the DMA operation has ended, software should poll the DMA_DONE bit, and not the DMA_RUN bit, to determine when the transaction has completed.

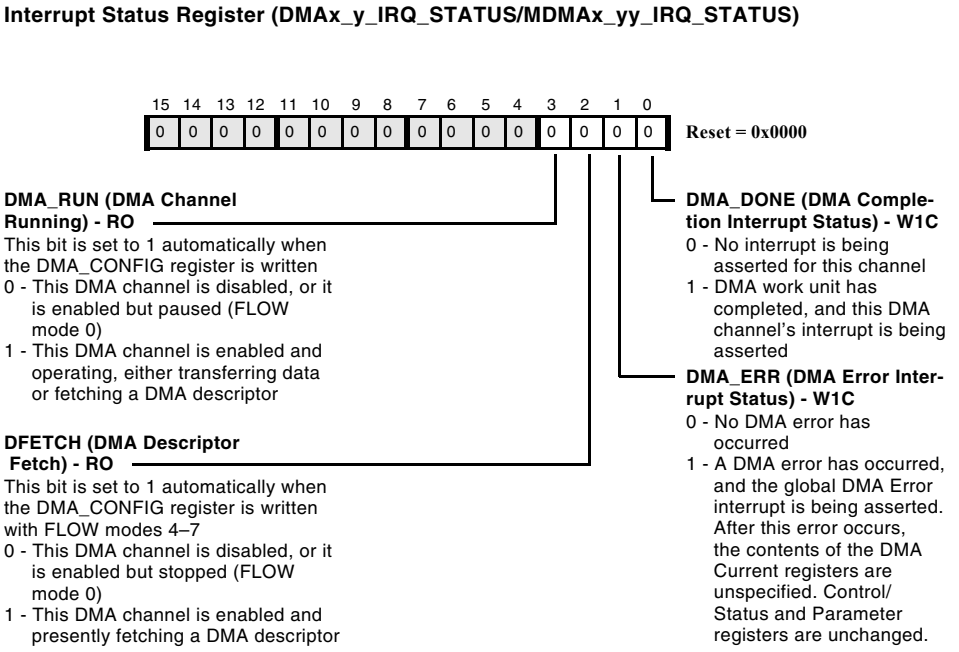

Figure 9-14. Interrupt Status Register

The processor supports a flexible interrupt control structure with three interrupt sources.

- Data driven interrupts (see Table 9-6)
- Peripheral Error interrupts
- DMA Error interrupts (for example, Bad Descriptor or Bus Error)

Separate Interrupt Request (IRQ) levels are allocated for Data and Peripheral Error interrupts, and DMA Error interrupts.

Table 9-6. Data Driven Interrupts

Interrupt Name	Description
No Interrupt	Interrupts can be disabled for a given work unit.
Peripheral Interrupt	These are peripheral (non-DMA) interrupts.
Row Completion	DMA Interrupts can occur on the completion of a row (CURR_X_COUNT expiration).
Buffer Completion	DMA Interrupts can occur on the completion of an entire buffer (when CURR_X_COUNT and CURR_Y_COUNT expire).

All DMA channels are ORed together into one system-level DMA Error interrupt. The individual IRQ_STATUS words of each channel can be read to identify the channel that caused the DMA Error interrupt.

Note the DMA_DONE and DMA_ERR interrupt indicators are write-one-to-clear (W1C).

When switching a peripheral from DMA to non-DMA mode, the peripheral's interrupts should be disabled during the mode switch (via the appropriate peripheral registers or SIC_IMASK) so that no unintended interrupt is generated on the shared DMA/interrupt request line.

Flex Descriptor Structure

DMA flex descriptors are variable sized data structures whose contents are loaded into DMA Parameter registers. The sequence of registers in the descriptor is essentially fixed (among three similar variations), but the length of the descriptor is completely programmable. The DMA channel registers are ordered so that the registers that are most commonly reloaded per work unit are at the lowest MMR addresses. The user may choose whether or not to use descriptors. If not using descriptors, the user can write the DMA MMRs directly to start DMA, and use either Autobuffer mode for continuous operation or Stop mode for single-buffer operation.

To use descriptors, the user programs the NDSIZE field of the DMAx_y_CONFIG register with the number of DMA registers to load from the descriptor, starting with the lowest MMR address. The user may select a descriptor size from one entry (the lower 16 bits of START_ADDR) to nine entries (all the DMA parameters).

The three variations of the descriptor value sequences depend on whether a Next Descriptor Pointer is included and, if so, what kind.

- None included (Descriptor Array mode)
- The lower 16 bits of the Next Descriptor Pointer (Descriptor List, Small Model)
- All 32 bits of the Next Descriptor Pointer (Descriptor List, Large Model)

All the other registers not loaded from the descriptor retain their prior values, although the CURR_ADDR, CURR_X_COUNT, and CURR_Y_COUNT registers are reloaded between the descriptor fetch and the start of DMA operation.

There are certain DMA settings that are not allowed to change from one descriptor to the next in a chain (Small or Large List and Array modes). These are DMA Direction, Word Size, and Memory Space (that is, switching between internal and external memory).

A single descriptor chain cannot control the transfer of a sequence of data buffers which reside in different memory spaces. Instead, group the data buffers into chains of buffers in the same space, but do not link the chains together. Transfer the first chain, wait for its final interrupt, and then start the next chain with an MMR write to the DMA_CONFIG register.

Note that while the user must locate each chain's data buffers in the same memory space, the descriptor structures themselves may be placed in any memory space, and they may link from a descriptor in one space to a descriptor in another space without restriction.

Table 9-7 shows the offsets for descriptor elements in the three modes described above. Note the names in the table list the descriptor elements in memory, not the actual MMRs into which they are eventually loaded.

Table 9-7. Parameter Registers and Descriptor Offsets

Descriptor Offset	Descriptor Array Mode	Small Descriptor List Mode	Large Descriptor List Mode
0x0	SAL	NDPL	NDPL
0x2	SAH	SAL	NDPH
0x4	DMACFG	SAH	SAL
0x6	XCNT	DMACFG	SAH
0x8	XMOD	XCNT	DMACFG
0xA	YCNT	XMOD	XCNT
0xC	YMOD	YCNT	XMOD
0xE		YMOD	YCNT
0x10			YMOD

DMA Operation Flow

Figure 9-15 and Figure 9-16 describe the DMA Flow.

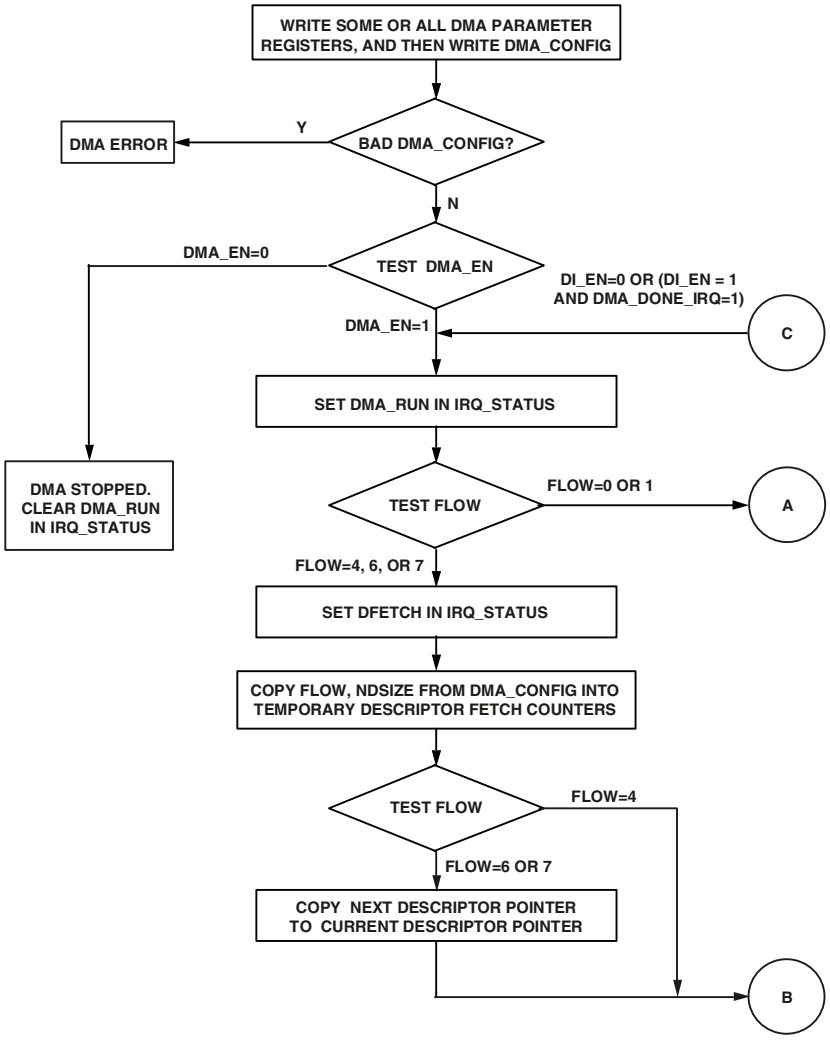

Figure 9-15. DMA Flow, From DMA Controller's Point of View (1 of 2)

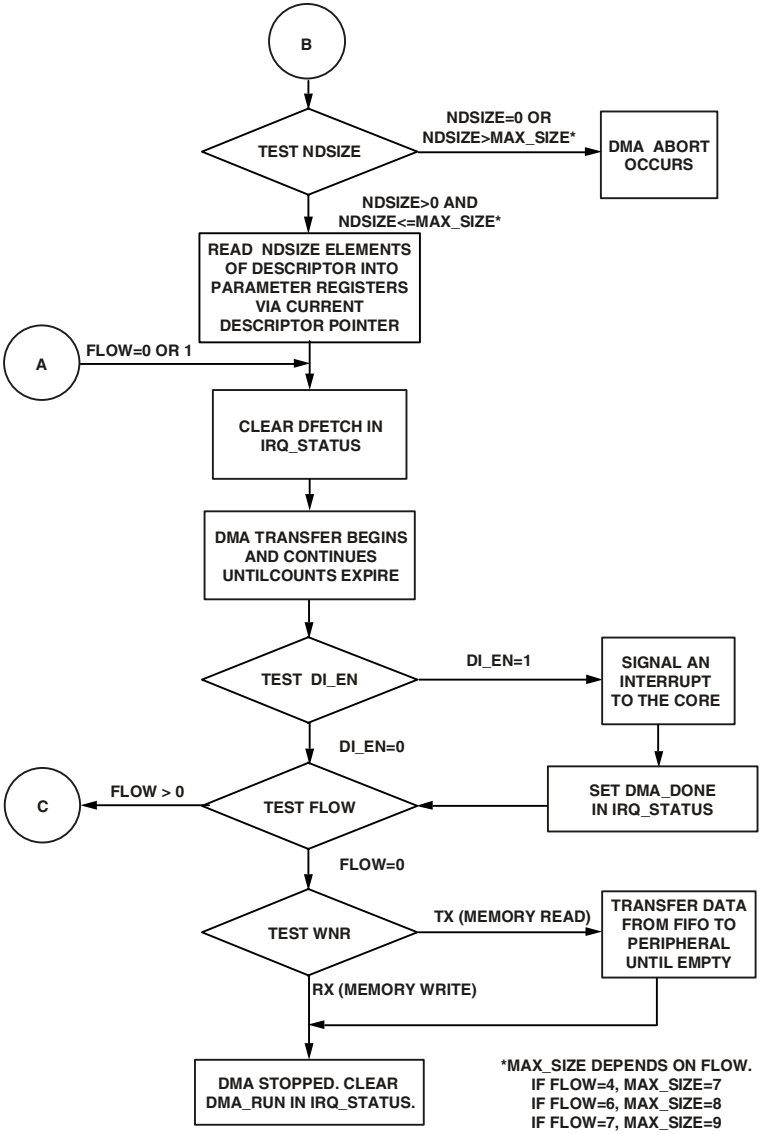

Figure 9-16. DMA Flow, From DMA Controller's Point of View (2 of 2)

DMA Startup

This section discusses starting DMA "from scratch." This is similar to starting it after it has been paused by FLOW = 0 mode.

Before initiating DMA for the first time on a given channel, be sure to initialize all Parameter registers. Be especially careful to initialize the upper 16 bits of the NEXT_DESC_PTR and START_ADDR registers, because they might not otherwise be accessed, depending on the chosen FLOW mode of operation.

To start DMA operation on a given channel, some or all of the DMA Parameter registers must first be written directly. At a minimum, the NEXT_DESC_PTR register (or CURR_DESC_PTR register in FLOW = 4 mode) must be written at this stage, but the user may wish to write other DMA registers that might be static throughout the course of DMA activity (for example, X_MODIFY, Y_MODIFY). The contents of NDSIZE and FLOW in DMA_CONFIG indicate which registers, if any, are fetched from descriptor elements in memory. After the descriptor fetch, if any, is completed, DMA operation begins, initiated by writing DMA_CONFIG with DMAEN = 1.

When DMA_CONFIG is written directly, the DMA controller recognizes this as the special startup condition that occurs when starting DMA for the first time on this channel or after the engine has been stopped (FLOW = 0).

When the descriptor fetch is complete and DMAEN = 1, the DMACFG descriptor element that was read into DMA_CONFIG assumes control. Before this point, the direct write to DMA_CONFIG had control. In other words, the WDSIZE, DI_EN, DI_SEL, RESTART, and DMA2D fields will be taken from the DMACFG value in the descriptor read from memory, while these field values initially written to the DMA_CONFIG register are ignored.

As Figure 9-15 and Figure 9-16 show, at startup the FLOW and NDSIZE bits in DMA_CONFIG determine the course of the DMA setup process. The FLOW value determines whether to load more Current registers from descriptor elements in memory, while the NDSIZE bits detail how many descriptor

elements to fetch before starting DMA. DMA registers not included in the descriptor are not modified from their prior values.

If the FLOW value specifies Small or Large Descriptor List modes, the NEXT_DESC_PTR is copied into CURR_DESC_PTR. Then, fetches of new descriptor elements from memory are performed, indexed by CURR_DESC_PTR, which is incremented after each fetch. If NDPL and/or NDPH is part of the descriptor, then these values are loaded into NEXT_DESC_PTR, but the fetch of the current descriptor continues using CURR_DESC_PTR. After completion of the descriptor fetch, CURR_DESC_PTR points to the next 16-bit word in memory past the end of the descriptor.

If neither NDPH nor NDPL are part of the descriptor (that is, in Descriptor Array mode, FLOW = 4), then the transfer from NDPH/NDPL into CURR_DESC_PTR does not occur. Instead, descriptor fetch indexing begins with the value in CURR_DESC_PTR.

If DMACFG is not part of the descriptor, the previous DMA_CONFIG settings (as written by MMR access at startup) control the work unit operation. If DMACFG is part of the descriptor, then the DMA_CONFIG value programmed by the MMR access controls only the loading of the first descriptor from memory. The subsequent DMA work operation is controlled by the low byte of the descriptor's DMACFG and by the Parameter registers loaded from the descriptor. The bits DI_EN, DI_SEL, DMA2D, WDSIZE, and WNR in the value programmed by the MMR access are disregarded.

The DMA_RUN and DFETCH status bits in the IRQ_STATUS register indicate the state of the DMA channel. After a write to DMA_CONFIG, the DMA_RUN and DFETCH bits can be automatically set to 1. No data interrupts are signaled as a result of loading the first descriptor from memory.

After the above steps, the Current registers are loaded automatically from the appropriate descriptor elements, overwriting their previous contents, as follows.

- START_ADDR is copied to CURR_ADDR
- X_COUNT is copied to CURR_X_COUNT
- Y_COUNT is copied to CURR_Y_COUNT

Then DMA data transfer operation begins, as shown in Figure 9-16.

DMA Refresh

On completion of a work unit, the DMA controller:

- Completes the transfer of all data between memory and the DMA unit.
- If enabled by DI_EN, signals an interrupt to the core and sets the DMA_DONE bit in the channel's IRQ_STATUS register.
- If FLOW = 0 (Stop) only:

  Stops operation by clearing the DMA_RUN bit in IRQ_STATUS after any data in the channel's DMA FIFO has been transferred to the peripheral.
- During the fetch in FLOW modes 4, 6, and 7, the DMA controller sets the DFETCH bit in IRQ_STATUS to 1. At this point, the DMA operation depends on whether FLOW = 4, 6, or 7, as follows:

  If FLOW = 4 (Descriptor Array):

  Loads a new descriptor from memory into DMA registers via the contents of CURR_DESC_PTR, while incrementing CURR_DESC_PTR. The descriptor size comes from the NDSIZE field of the DMA_CONFIG value prior to the beginning of the fetch.

If FLOW = 6 (Descriptor List Small):

Copies the 32-bit NEXT_DESC_PTR into CURR_DESC_PTR. Next, fetches a descriptor from memory into DMA registers via the new contents of CURR_DESC_PTR, while incrementing CURR_DESC_PTR. The first descriptor element loaded is a new 16-bit value for the lower 16 bits of NEXT_DESC_PTR, followed by the rest of the descriptor elements. The high 16 bits of NEXT_DESC_PTR will retain their former value. This supports a shorter, more efficient descriptor than the Descriptor List Large model, suitable whenever the application can place the channel's descriptors in the same 64K byte range of memory.

If FLOW = 7 (Descriptor List Large):

Copies the 32-bit NEXT_DESC_PTR into CURR_DESC_PTR. Next, fetches a descriptor from memory into DMA registers via the new contents of CURR_DESC_PTR, while incrementing CURR_DESC_PTR. The first descriptor element loaded is a new 32-bit value for the full NEXT_DESC_PTR, followed by the rest of the descriptor elements. The high 16 bits of NEXT_DESC_PTR may differ from their former value. This supports a fully flexible descriptor list which can be located anywhere in internal memory, or external memory.

Note if it is necessary to link from a descriptor chain whose descriptors are in one 64K byte area to another chain whose descriptors are outside that area, only one descriptor needs to use FLOW = 7—just the descriptor which contains the link leaving the 64K byte range. All the other descriptors located together in the same 64K byte areas may use FLOW = 6.

- If FLOW = 1, 4, 6, or 7 (Autobuffer, Descriptor Array, Descriptor List Small, or Descriptor List Large, respectively):

  (Re)loads the Current registers:

  CURR_ADDR loaded from START_ADDR

  CURR_X_COUNT loaded from X_COUNT

  CURR_Y_COUNT loaded from Y_COUNT

  The DFETCH bit in IRQ_STATUS is then cleared, after which the DMA transfer begins again, as shown in Figure 9-16.

To Stop DMA Transfers

In FLOW = 0 mode, DMA stops automatically after the work unit is complete.

If a list or array of descriptors is used to control DMA, and if every descriptor contains a DMACFG element, then the final DMACFG element should have a FLOW = 0 setting to gracefully stop the channel.

In Autobuffer (FLOW = 1) mode, or if a list or array of descriptors without DMACFG elements is used, then the DMA transfer process must be terminated by an MMR write to the DMAx_y_CONFIG register with a value whose DMAEN bit is 0. A write of 0 to the entire register will always terminate DMA gracefully (without DMA Abort).

Before enabling the channel again, make sure that any slow memory read operations that may have started are completed (for example, reads from slow external memory). Do not enable the channel again until any such reads are complete.

Two-Dimensional DMA

Two-dimensional (2D) DMA supports arbitrary row and column sizes up to 64 K × 64 K elements, as well as arbitrary X_MODIFY and Y_MODIFY values up to ±32 K bytes. Furthermore, Y_MODIFY can be negative, allowing implementation of interleaved datastreams. The X_COUNT and Y_COUNT values specify the row and column sizes, where X_COUNT must be 2 or greater.

The start address and modify values are in bytes, and they must be aligned to a multiple of the DMA transfer word size (WDSIZE[1:0] in DMA_CONFIG). Misalignment causes a DMA error.

The X_MODIFY value is the byte-address increment that is applied after each transfer that decrements the CURR_X_COUNT register. The X_MODIFY value is not applied when the inner loop count is ended by decrementing CURR_X_COUNT from 1 to 0, except that it is applied on the final transfer when CURR_Y_COUNT is 1 and CURR_X_COUNT decrements from 1 to 0.

The Y_MODIFY value is the byte-address increment that is applied after each decrement of CURR_Y_COUNT. However, the Y_MODIFY value is not applied to the last item in the array on which the outer loop count (CURR_Y_COUNT) also expires by decrementing from 1 to 0.

After the last transfer completes, CURR_Y_COUNT = 1, CURR_X_COUNT = 0, and CURR_ADDR is equal to the last item's address plus X_MODIFY. Note if the DMA channel is programmed to refresh automatically (Autobuffer mode), then these registers will be loaded from X_COUNT, Y_COUNT, and START_ADDR upon the first data transfer.

Example 1: Retrieve a 16 × 8 block of bytes from a video frame buffer of size (N × M) pixels:

```
X_MODIFY = 1
X_COUNT = 16
```

```
Y_MODIFY = N-15 (offset from the end of one row to the start of
another)
Y_COUNT = 8
```

This produces the following address offsets from the start address:

```
0,1,2,...15,
N,N + 1, ... N + 15,
2N, 2N + 1,... 2N + 15, ...
7N, 7N + 1,... 7N + 15,
```

Example 2: Receive a video datastream of bytes,
(R,G,B pixels) × (N × M image size):

```
X_MODIFY = (N * M)
X_COUNT = 3
Y_MODIFY = 1 - 2(N * M) (negative)
Y_COUNT = (N * M)
```

This produces the following address offsets from the start address:

```
0, (N * M), 2(N * M),
1, (N * M) + 1, 2(N * M) + 1,
2, (N * M) + 2, 2(N * M) + 2,
    ...
(N * M) - 1, 2(N * M) - 1, 3(N * M) - 1,
```

More 2D DMA Examples

Examples of DMA styles supported by flex descriptors include:

- A single linear buffer that stops on completion (`FLOW` = Stop mode)
- A linear buffer with stride greater than 1 (`X_MODIFY > 1`)
- A circular, auto-refreshing buffer that interrupts on each full buffer

- A similar buffer that interrupts on fractional buffers (for example, 1/2, 1/4) (2D DMA)
- 1D DMA, using a set of identical ping-pong buffers defined by a linked ring of 3-word descriptors, each containing {link pointer, 32-bit address}
- 1D DMA, using a linked list of 5-word descriptors containing { link pointer, 32-bit address, length, config}
- 2D DMA, using an array of 1-word descriptors, specifying only the base DMA address within a common data page
- 2D DMA, using a linked list of 9-word descriptors, specifying everything

Memory DMA

This section describes the Memory DMA (MDMAs, IMDMA) controller, which provides memory-to-memory DMA transfers among the processor's memory spaces. These include internal L1 and L2 memories and external synchronous/ asynchronous memories. IMDMA only supports data transfer within internal memories—for example, between L1 and L1, or L1 and L2, or L2 and L2).

Each Memory DMA controller contains a DMA FIFO, an 8-word by 32-bit FIFO block used to transfer data to and from either internal memory or the EAB bus. It will also support DMA from Boot ROM on the EAB bus. The FIFO can be used to hold DMA data transferred between two internal L1 memory locations or between two external memory locations.

The processor provides twelve channels for Memory DMA controller:

- Six source channels (for reading from memory)
- Six destination channels (for writing to memory)

Each source/destination channel forms a stream, and these two streams are fixed at DMA priorities 12 through 15 of DMA1 and DMA2, or 0 through 3 of IMDMA. The channels support 8-bit, 16-bit, and 32-bit Memory DMA transfers, but both ends of the Memory DMA transfer must be programmed to the same word size. In other words, the Memory DMA transfer does not perform packing or unpacking of data; each read results in one write. Both ends of Memory DMA FIFO for a given stream are granted priority at the same time. Each stream shares an 8-word-deep 32-bit FIFO. The source DMA engine fills the FIFO, while the destination DMA engine empties it. The FIFO depth allows the burst transfers of the EAB and DAB buses to overlap, significantly improving throughput on block transfers between internal and external memory.

Two separate descriptor blocks are required to supply the operating parameters for each Memory DMA stream, one for the source channel and one for the destination channel.

Because the source and destination DMA engines share a single FIFO buffer, the descriptor blocks must be configured to have the same data size. It is possible to have a different mix of descriptors on both ends as long as the total count is the same.

To start an MDMA transfer operation, the MMRs for the source and destination streams are written, each in a manner similar to peripheral DMA.

The only constraint is that the DMA_CONFIG register for the source stream must be written before the DMA_CONFIG register for the destination stream.

When the destination DMA_CONFIG register is written, MDMA operation starts, after a latency of 3 SCLK cycles.

First, if either MDMA stream has been selected to use descriptors, the descriptors are fetched from memory. The destination stream descriptors are fetched first. Then, after a latency of 4 SCLK cycles after the last descriptor word is returned from memory (or typically 8 SCLK cycles after the fetch of the last descriptor word, due to memory pipelining), the source MDMA stream begins fetching data from the source buffer. The resulting data is deposited in the MDMA stream's 8-location FIFO, and then after a latency of 2 SCLK cycles, the destination MDMA stream begins writing data to the destination memory buffer.

MDMA Bandwidth

If source and destination are in different memory spaces (one internal and one external), the internal and external memory transfers are typically simultaneous and continuous, maintaining 100% bus utilization of the internal and external memory interfaces. This performance is affected by core-to-system clock frequency ratios. At ratios below about 2.5:1, synchronization and pipeline latencies result in lower bus utilization in the system clock domain. At a clock ratio of 2:1, for example, DMA typically runs at 2/3 of the system clock rate. At higher clock ratios, full bandwidth is maintained.

If source and destination are in the same memory space (both internal or both external), the MDMA stream typically prefetches a burst of source data into the FIFO, and then automatically turns around and delivers all available data from the FIFO to the destination buffer. The burst length is dependent on traffic, and is equal to 3 plus the memory latency at the DMA in SCLKs (which is typically 7 for internal transfers and 6 for external transfers).

DMA Performance Optimization

The DMA system is designed to provide maximum throughput per channel and maximum utilization of the internal buses, while accommodating the inherent latencies of memory accesses.

A key feature of the DMA architecture is the separation of the activity on the peripheral DMA bus (the DMA Access Bus (DAB)) from the activity on the buses between the DMA and memory (the DMA Core Bus (DCB) and the DMA External Bus (DEB)). Each peripheral DMA channel has its own data FIFO which lies between the DAB bus and the memory buses. These FIFOs automatically prefetch data from memory for transmission and buffer received data for later memory writes. This allows the peripheral to be granted a DMA transfer with very low latency compared to the total latency of a pipelined memory access, permitting the repeat rate (bandwidth) of each DMA channel to be as fast as possible.

Peripheral DMA channels have a maximum transfer rate of one 32/16-bit word (DMA1/DMA2) per two system clocks, per channel, in either direction.

MDMA channels have a maximum transfer rate of one 32-bit word per one system clock (`SCLK`), per channel.

When all DMA channels' traffic is taken in the aggregate:

- Transfers between the peripherals and the DMA unit have a maximum rate of one 32/16-bit transfer (DMA1/DMA2) per system clock.
- Transfers between the DMA unit and internal memory (L1) have a maximum rate of one 32-bit transfer per system clock.
- Transfers between the DMA unit and external memory have a maximum rate of one 32-bit transfer per system clock.

Some considerations which limit the actual performance include:

- Accesses to internal or external memory which conflict with core accesses to the same memory. This can cause delays, for example, for accessing the same L1 bank, for opening/closing SDRAM pages, or while filling cache lines.
- Each direction change from RX to TX on the DAB bus imposes a one SCLK cycle delay.
- Direction changes on the DCB bus (for example, write followed by read) to the same bank of internal memory can impose delays.
- Direction changes (for example, read followed by write) on the DEB bus to external memory can each impose a several-cycle delay.
- MMR accesses to DMA registers other than DMAx_y_CONFIG, DMAx_y_IRQSTAT, or DMAx_y_PERIPHERAL_MAP will stall all DMA activity for one cycle per 32-bit word transferred. In contrast, MMR accesses to the Control/Status registers do not cause stalls or wait states.
- Reads from DMA registers other than Control/Status registers use one PAB bus wait state, delaying the core for several core clocks.
- Descriptor fetches consume one DMA memory cycle per 32-bit word read from memory, but do not delay transfers on the DAB bus.
- Initialization of a DMA channel stalls DMA activity for one cycle. This occurs when DMAEN changes from 0 to 1 or when the RESTART bit is set to 1 in the DMAx_y_CONFIG register.

Several of these factors may be minimized by proper design of the application software. It is often possible to structure the software to avoid internal and external memory conflicts by careful allocation of data buffers within banks and pages, and by planning for low cache activity during

critical DMA operations. Furthermore, unnecessary MMR accesses can be minimized, especially by using descriptors or autobuffering.

Efficiency loss caused by excessive direction changes (thrashing) can be minimized by the processor's traffic control features, described in the next section.

Prioritization and Traffic Control

DMA channels are ordinarily granted service strictly according to their priority. The priority of a channel is simply its channel number, where lower priority numbers are granted first. Thus, peripherals with high data rates or low latency requirements should be assigned to lower numbered (higher priority) channels using the DMAx_y_PERIPHERAL_MAP registers. The Memory DMA streams are always lower priority than the peripherals, but as they request service continuously, they ensure that any time slots unused by peripheral DMA are applied to MDMA transfers. By default, when more than one MDMA stream is enabled and ready, only the highest priority MDMA stream is granted. If it is desirable for the MDMA streams to share the available bandwidth, however, the MDMA_ROUND_ROBIN_PERIOD may be programmed to select each stream in turn for a fixed number of transfers.

In the processor DMA, there are two completely separate but simultaneous prioritization processes—the DAB bus prioritization and the memory bus (DCB and DEB) prioritization. Peripherals that are requesting DMA via the DAB bus, and whose data FIFOs are ready to handle the transfer, compete with each other for DAB bus cycles. Similarly but separately, channels whose FIFOs need memory service (prefetch or post-write) compete together for access to the memory buses. MDMA streams compete for memory access as a unit, and source and destination may be granted together if their memory transfers do not conflict. In this way, internal-to-external or external-to-internal memory transfers may occur at the full system clock rate (SCLK). Examples of memory conflict include simultaneous access to the same memory space and simultaneous

attempts to fetch descriptors. Special processing may occur if a peripheral is requesting DMA but its FIFO is not ready (for example, an empty transmit FIFO or full receive FIFO). For more information, see "Urgent DMA Transfers" on page 9-51.

If a descriptor fetch is in progress for a DMA channel (with descriptor in on-chip or off-chip/external memory)—another channel descriptor fetch (with descriptor in the other memory) cannot occur until the first channel is finished, even if it has higher priority.

Traffic control is an important consideration in optimizing use of DMA resources. Traffic control is a way to influence how often the transfer direction on the data buses may change, by automatically grouping same direction transfers together. Each DMA block provides a traffic control mechanism controlled by the `DMAx_TC_PER` and `DMAx_TC_CNT` registers. This mechanism performs the optimization without real-time processor intervention, and without the need to program transfer bursts into the DMA work unit streams. Traffic can be independently controlled for each of the three buses (DAB, DCB, and DEB) with simple counters. In addition, alternation of transfers among MDMA streams can be controlled with the `MDMA_ROUND_ROBIN_COUNT` field of the `DMA_TC_CNT` register. See "MDMA Priority and Scheduling" on page 9-50.

Using the traffic control features, the DMA system preferentially grants data transfers on the DAB or memory buses which are going in the same read/write direction as the previous transfer, until either the traffic control counter times out, or until traffic stops or changes direction on its own. When the traffic counter reaches zero, the preference is changed to the opposite flow direction. These directional preferences work as if the priority of the opposite direction channels were decreased by 16.

For example, if channels 3 and 5 were requesting DAB access, but lower priority channel 5 is going "with traffic" and higher priority channel 3 is going "against traffic," then channel 3's effective priority becomes 19, and channel 5 would be granted instead. If, on the next cycle, only channels 3

and 6 were requesting DAB transfers, and these transfer requests were both “against traffic,” then their effective priorities would become 19 and 22, respectively. One of the channels (channel 3) is granted, even though its direction is opposite to the current flow. No bus cycles are wasted, other than any necessary delay required by the bus turnaround.

This type of traffic control represents a trade-off of latency to improve utilization (efficiency). Higher traffic timeouts might increase the length of time each request waits for its grant, but it often dramatically improves the maximum attainable bandwidth in congested systems, often to above 90%.

To disable preferential DMA prioritization, program the appropriate DMAx_TC_PER register to 0x0000. Each of the system DMA controllers (DMA1 and DMA2) has a set of DMAx_TC_PER and DMAx_TC_CNT registers.

DMA Traffic Control Counter Period Register (DMAx_TC_PER) and DMA Traffic Control Counter Register (DMAx_TC_CNT)

The MDMA_ROUND_ROBIN_COUNT field shows the current transfer count remaining in the MDMA round robin period. It initializes to MDMA_ROUND_ROBIN_PERIOD whenever DMAx_TC_PER is written, whenever a different MDMA stream is granted, or whenever every MDMA stream is idle. It then counts down to 0 with each MDMA transfer. When this count decrements from 1 to 0, the next available MDMA stream is selected.

The DAB_TRAFFIC_COUNT field shows the current cycle count remaining in the DAB traffic period. It initializes to DAB_TRAFFIC_PERIOD whenever DMAx_TC_PER is written, or whenever the DAB bus changes direction or becomes idle. It then counts down from DAB_TRAFFIC_PERIOD to 0 on each system clock (except for DMA stalls). While this count is nonzero, same direction DAB accesses are treated preferentially. When this count decrements from 1 to 0, the opposite direction DAB access is treated preferentially, which may result in a direction change. When this count is

0 and a DAB bus access occurs, the count is reloaded from DAB_TRAFFIC_PERIOD to begin a new burst.

The DEB_TRAFFIC_COUNT field shows the current cycle count remaining in the DEB traffic period. It initializes to DEB_TRAFFIC_PERIOD whenever DMAx_TC_PER is written, or whenever the DEB bus changes direction or becomes idle. It then counts down from DEB_TRAFFIC_PERIOD to 0 on each system clock (except for DMA stalls). While this count is nonzero, same direction DEB accesses are treated preferentially. When this count decrements from 1 to 0, the opposite direction DEB access is treated preferentially, which may result in a direction change. When this count is 0 and a DEB bus access occurs, the count is reloaded from DEB_TRAFFIC_PERIOD to begin a new burst.

The DCB_TRAFFIC_COUNT field shows the current cycle count remaining in the DCB traffic period. It initializes to DCB_TRAFFIC_PERIOD whenever DMAx_TC_PER is written, or whenever the DCB bus changes direction or becomes idle. It then counts down from DCB_TRAFFIC_PERIOD to 0 on each system clock (except for DMA stalls). While this count is nonzero, same direction DCB accesses are treated preferentially. When this count decrements from 1 to 0, the opposite direction DCB access is treated preferentially, which may result in a direction change. When this count is 0 and a DCB bus access occurs, the count is reloaded from DCB_TRAFFIC_PERIOD to begin a new burst. See Figure 9-17 and Figure 9-18.

DMA Traffic Control Counter Period Register (DMAx_TC_PER)

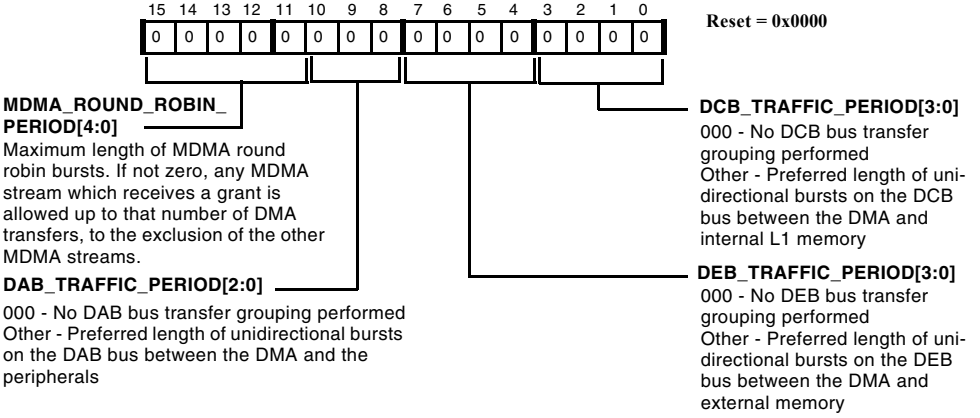

Figure 9-17. DMA Traffic Control Counter Period Register

DMA Traffic Control Counter Register (DMAx_TC_CNT)
RO

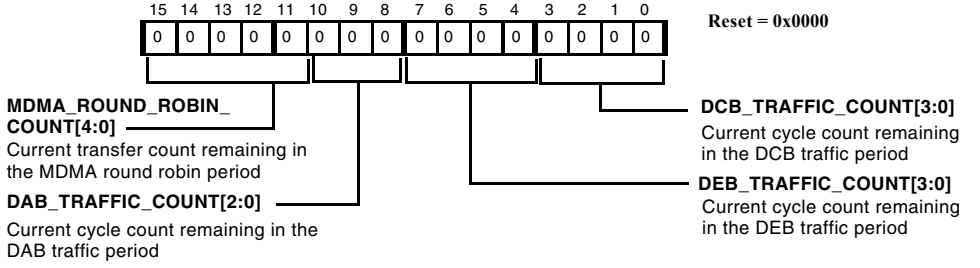

Figure 9-18. DMA Traffic Control Counter Register

MDMA Priority and Scheduling

All MDMA operations have lower precedence than any peripheral DMA operations. MDMA thus makes effective use of any memory bandwidth unused by peripheral DMA traffic.

If two MDMA streams are used (S0-D0 and S1-D1), the user may choose to allocate bandwidth either by fixed stream priority or by a round robin scheme. This is selected by programming the `MDMA_ROUND_ROBIN_PERIOD` field in the `DMAx_TC_PER` register (see "Prioritization and Traffic Control" on page 9-45).

If this field is set to 0, then MDMA is scheduled by fixed priority. MDMA Stream 0 takes precedence over MDMA Stream 1 whenever Stream 0 is ready to perform transfers. Since an MDMA Stream is typically capable of transferring data on every available cycle, this could cause MDMA Stream 1 traffic to be delayed for an indefinite time until any and all MDMA Stream 0 operations are complete. This scheme could be appropriate in systems where low duration but latency sensitive data buffers need to be moved immediately, interrupting long duration, low priority background transfers.

If the `MDMA_ROUND_ROBIN_PERIOD` field is set to some nonzero value in the range `1 <= P <= 31`, then a round robin scheduling method is used. The two MDMA streams are granted bus access in alternation in bursts of up to `P` data transfers. This could be used in systems where two transfer processes need to coexist, each with a guaranteed fraction of the available bandwidth. For example, one stream might be programmed for internal-to-external moves while the other is programmed for external-to-internal moves, and each would be allocated approximately equal data bandwidth.

In round robin operation, the MDMA stream selection at any time is either "free" or "locked." Initially, the selection is free. On any free cycle available to MDMA (when no peripheral DMA accesses take precedence), if either or both MDMA streams request access, the higher precedence

stream will be granted (Stream 0 in case of conflict), and that stream's selection is then "locked." The MDMA_ROUND_ROBIN_COUNT counter field in the DMAx_TC_CNT register is loaded with the period P from MDMA_ROUND_ROBIN_PERIOD, and MDMA transfers begin. The counter is decremented on every data transfer (as each data word is written to memory). After the transfer corresponding to a count of 1, the MDMA stream selection is passed automatically to the other stream with zero overhead, and the MDMA_ROUND_ROBIN_COUNT counter is reloaded with the period value P from MDMA_ROUND_ROBIN_PERIOD. In this cycle, if the other MDMA stream is ready to perform a transfer, the stream selection is locked on the new MDMA stream. If the other MDMA stream is not ready to perform a transfer, then no transfer is performed, and on the next cycle the stream selection unlocks and becomes free again.

If round robin operation is used when only one MDMA stream is active, one idle cycle will occur for each P MDMA data cycles, slightly lowering bandwidth by a factor of 1/(P+1). If both MDMA streams are used, however, memory DMA can operate continuously with zero additional overhead for alternation of streams (other than overhead cycles normally associated with reversal of read/write direction to memory, for example). By selection of various round robin period values P which limit how often the MDMA streams alternate, maximal transfer efficiency can be maintained.

Urgent DMA Transfers

Typically, DMA transfers for a given peripheral occur at regular intervals. Generally, the shorter the interval, the higher the priority that should be assigned to the peripheral. If the average bandwidth of all the peripherals is not too large a fraction of the total, then all peripherals' requests should be granted as required.

Occasionally, instantaneous DMA traffic might exceed the available bandwidth, causing congestion. This may occur if L1 or external memory is temporarily stalled, perhaps for an SDRAM page swap or a cache line fill.

Congestion might also occur if one or more DMA channels initiates a flurry of requests, perhaps for descriptor fetches or to fill a FIFO in the DMA or in the peripheral.

If congestion persists, lower priority DMA peripherals may become starved for data. Even though the peripheral's priority is low, if the necessary data transfer does not take place before the end of the peripheral's regular interval, system failure may result. To minimize this possibility, the DMA unit detects peripherals whose need for data has become urgent, and preferentially grants them service at the highest priority.

A DMA channel's request for memory service is defined as "urgent" if both:

- The channel's FIFO is not ready for a DAB bus transfer (that is, a transmit FIFO is empty or a receive FIFO is full), and
- The peripheral is asserting its DMA request line.

Descriptor fetches may be urgent, if they are necessary to initiate or continue a DMA work unit chain for a starving peripheral. DMA requests from an MDMA channel are never urgent.

When one or more DMA channels express an urgent memory request, two events occur:

- All non-urgent memory requests are decreased in priority by 32, guaranteeing that only an urgent request will be granted. The urgent requests compete with each other, if there is more than one, and directional preference among urgent requests is observed.
- The resulting memory transfer is marked for expedited processing in the targeted memory system (L1 or external), and so are all prior incomplete memory transfers ahead of it in that memory system. This may cause a series of external memory core accesses to be delayed for a few cycles so that a peripheral's urgent request may be accommodated.

The preferential handling of urgent DMA transfers is completely automatic. No user controls are required for this function to operate.

Software Management of DMA

Several synchronization and control methods are available for use in development of software tasks which manage DMA and MDMA (see also "Memory DMA" on page 9-40). Such software needs to be able to accept requests for new DMA transfers from other software tasks, integrate these transfers into existing transfer queues, and reliably notify other tasks when the transfers are complete.

In the processor, it is possible for each DMA peripheral and MDMA stream to be managed by a separate task or to be managed together with any other stream. Each DMA channel has independent, orthogonal control registers, resources, and interrupts, so that the selection of the control scheme for one channel does not affect the choice of control scheme on other channels. For example, one peripheral can use a linked-descriptor-list, interrupt-driven scheme while another peripheral can simultaneously use a demand-driven, buffer-at-a-time scheme synchronized by polling of the `IRQ_STATUS` register.

Synchronization of Software and DMA

A critical element of software DMA management is synchronization of DMA buffer completion with the software. This can best be done using interrupts, polling of `IRQ_STATUS`, or a combination of both. Polling for address or count can only provide synchronization within loose tolerances comparable to pipeline lengths.

Interrupt-based synchronization methods must avoid interrupt overrun, or the failure to invoke a DMA channel's interrupt handler for every interrupt event due to excessive latency in processing of interrupts. Generally, the system design must either ensure that only one interrupt per channel is

scheduled (for example, at the end of a descriptor list), or that interrupts are spaced sufficiently far apart in time that system processing budgets can guarantee every interrupt is serviced. Note since every interrupt channel has its own distinct interrupt, interaction among the interrupts of different peripherals is much simpler to manage.

Polling of the CURR_ADDR, CURR_DESC_PTR, or CURR_X/Y_COUNT registers is not recommended as a method of precisely synchronizing DMA with data processing, due to DMA FIFOs and DMA/memory pipelining. The Current Address, Pointer, and Count registers change several cycles in advance of the completion of the corresponding memory operation, as measured by the time at which the results of the operation would first be visible to the core by memory read or write instructions. For example, in a DMA memory write operation to external memory, assume a DMA write by channel A is initiated that causes the SDRAM to perform a page open operation which will take many system clock cycles. The DMA engine may then move on to another DMA operation by channel B which does not in itself incur latency, but will be stalled behind the slow operation by channel A. Software monitoring channel B could not safely conclude whether the memory location pointed to by channel B's CURR_ADDR has or has not been written, based on examination of the CURR_ADDR register contents.

Polling of the Current Address, Pointer, and Count registers can permit loose synchronization of DMA with software, however, if allowances are made for the lengths of the DMA/memory pipeline. The length of the DMA FIFO for a peripheral DMA channel is four locations (either four 8- or 16-bit data elements, or two 32-bit data elements) and for an MDMA FIFO is eight locations (four 32-bit data elements). The DMA will not advance Current Address/Pointer/Count registers if these FIFOs are filled with incomplete work (including reads that have been started but not yet finished).

Additionally, the length of the combined DMA and L1 pipelines to internal memory is approximately six 8- or 16-bit data elements. The length of

the DMA and External Bus Interface Unit (EBIU) pipelines is approximately three data elements, when measured from the point where a DMA register update is visible to an MMR read to the point where DMA and core accesses to memory become strictly ordered. If the DMA FIFO length and the DMA/memory pipeline length are added, an estimate can be made of the maximum number of incomplete memory operations in progress at one time. (Note this is a maximum, as the DMA/memory pipeline may include traffic from other DMA channels.)

For example, assume a peripheral DMA channel is transferring a work unit of 100 data elements into internal memory and its `CURR_X_COUNT` register reads a value of 60 remaining elements, so that processing of the first 40 elements has at least been started. The total pipeline length is no greater than the sum of 4 (for the PDMA FIFO) plus 6 (for the DMA/memory pipeline), or 10 data elements, so it is safe to conclude that the DMA transfer of the first 40-10 = 30 data elements is complete.

For precise synchronization, software should either wait for an interrupt or consult the channel's `IRQ_STATUS` register to confirm completion of DMA, rather than polling Current Address/Pointer/Count registers. When the DMA system issues an interrupt or changes an `IRQ_STATUS` bit, it guarantees that the last memory operation of the work unit has been completed and will definitely be visible to DSP code. For memory read DMA, the final memory read data will have been safely received in the DMA's FIFO; for memory write DMA, the DMA unit will have received an acknowledge from L1 memory or the EBIU that the data has been written.

The following examples show methods of synchronizing software with several different styles of DMA.

Single-Buffer DMA Transfers

Synchronization is simple if a peripheral's DMA activity consists of isolated transfers of single buffers. DMA activity is initiated by software writes to the channel's Control registers. The user may choose to use a

single descriptor in memory, in which case the software only needs to write the DMA_CONFIG and the NEXT_DESC_PTR registers. Alternatively, the user may choose to write all the MMR registers directly from software, ending with the write to the DMA_CONFIG register.

The simplest way to signal completion of DMA is by an interrupt. This is selected by the DI_EN bit in the DMA_CONFIG register, and by the necessary setup of the System Interrupt Controller. If it is desirable not to use an interrupt, the software can poll for completion by reading the IRQ_STATUS register and testing the DMA_RUN bit. If this bit is zero, the buffer transfer has completed.

Continuous Transfers Using Autobuffering

If a peripheral's DMA data consists of a steady, periodic stream of signal data, DMA autobuffering (FLOW = 1) may be an effective option. Here, DMA is transferred from or to a memory buffer with a circular addressing scheme, using either one- or two-dimensional indexing with zero processor and DMA overhead for looping. Synchronization options include:

- 1D, interrupt-driven—software is interrupted at the conclusion of each buffer. The critical design consideration is that the software must deal with the first items in the buffer before the next DMA transfer, which might overwrite or re-read the first buffer location before it is processed by software. This scheme may be workable if the system design guarantees that the data repeat period is longer than the interrupt latency under all circumstances.
- 2D, interrupt-driven (double buffering)—the DMA buffer is partitioned into two or more sub-buffers, and interrupts are selected (set DI_SEL = 1 in DMA_CONFIG) to be signaled at the completion of

each DMA inner loop. For example, two 512-word sub-buffers inside a 1K word buffer could be used to receive 16-bit peripheral data with these settings:

- `START_ADDR` = buffer base address
- `DMA_CONFIG = 0x10D7` (`FLOW = 1`, `DI_EN = 1`, `DI_SEL = 1`, `DMA2D = 1`, `WDSIZE = 01`, `WNR = 1`, `DMAEN = 1`)
- `X_COUNT = 512`
- `X_MODIFY = 2` for 16-bit data
- `Y_COUNT = 2` for two sub-buffers
- `Y_MODIFY = 2`, same as `X_MODIFY` for contiguous sub-buffers

In this way, a traditional double buffer or “ping-pong” scheme could be implemented.

- 2D, polled—if interrupt overhead is unacceptable but the loose synchronization of address/count register polling is acceptable, a 2D multibuffer synchronization scheme may be used. For example, assume receive data needs to be processed in packets of sixteen 32-bit elements. A four-part 2D DMA buffer can be allocated where each of the four sub-buffers can hold one packet with these settings:
  - `START_ADDR` = buffer base address
  - `DMA_CONFIG = 0x101B` (`FLOW = 1`, `DI_EN = 0`, `DMA2D = 1`, `WDSIZE = 10`, `WNR = 1`, `DMAEN = 1`)
  - `X_COUNT = 16`
  - `X_MODIFY = 4` for 32-bit data

- `Y_COUNT = 4` for four sub-buffers
- `Y_MODIFY = 4`, same as `X_MODIFY` for contiguous sub-buffers

The synchronization core might read `Y_COUNT` to determine which sub-buffer is currently being transferred, and then allow one full sub-buffer to account for pipelining. For example, if a read of `Y_COUNT` shows a value of 3, then the software should assume that sub-buffer 3 is being transferred, but some portion of sub-buffer 2 may not yet be received. The software could, however, safely proceed with processing sub-buffers 1 or 0.

- 1D unsynchronized FIFO—if a system's design guarantees that the processing of a peripheral's data and the DMA rate of the data will remain correlated in the steady state, but that short-term latency variations must be tolerated, it may be appropriate to build a simple FIFO. Here, the DMA channel may be programmed using 1D Autobuffer mode addressing without any interrupts or polling.

Descriptor Structures

DMA descriptors may be used to transfer data to or from memory data structures that are not simple 1D or 2D arrays. For example, if a packet of data is to be transmitted from several different locations in memory (a header from one location, a payload from a list of several blocks of memory managed by a memory pool allocator, and a small trailer containing a checksum), a separate DMA descriptor can be prepared for each memory area, and the descriptors can be grouped in either an array or list as desired by selecting the appropriate `FLOW` setting in `DMA_CONFIG`.

The software can synchronize with the progress of the structure's transfer by selecting interrupt notification for one or more of the descriptors. For example, the software might select interrupt notification for the header's descriptor and for the trailer's descriptor, but not for the payload blocks' descriptors.

It is important to remember the meaning of the various fields in the DMA_CONFIG descriptor elements when building a list or array of DMA descriptors. In particular:

- The lower byte of DMA_CONFIG specifies the DMA transfer to be performed by the *current* descriptor (for example, interrupt-enable, 2D mode)
- The upper byte of DMA_CONFIG specifies the format of the *next* descriptor in the chain. The NDSIZE and FLOW fields in a given descriptor do not correspond to the format of the descriptor itself; they specify the link to the next descriptor, if any.

On the other hand, when the DMA unit is being restarted, both bytes of the DMA_CONFIG value written to the DMA channel's DMA_CONFIG register should correspond to the current descriptor. At a minimum, the FLOW, NDSIZE, WNR, and DMAEN fields must all agree with the current descriptor; the WDSIZE, DI_EN, DI_SEL, RESTART, and DMA2D fields will be taken from the DMA_CONFIG value in the descriptor read from memory (and the field values initially written to the register are ignored).

Descriptor Queue Management

A system designer might want to write a DMA Manager facility which accepts DMA requests from other software. The DMA Manager software does not know in advance when new work requests will be received or what these requests might contain. The software could manage these transfers using a circular linked list of DMA descriptors, where each descriptor's NDPH and NDPL members point to the next descriptor, and the last descriptor points to the first.

The code that writes into this descriptor list could use the processor's circular addressing modes (I, L, M, and B registers), so that it does not need to use comparison and conditional instructions to manage the circular structure. In this case, the NDPH and NDPL members of each descriptor could

even be written once at startup, and skipped over as each descriptor's new contents are written.

The recommended method for synchronization of a descriptor queue is through the use of an interrupt. The descriptor queue is structured so that at least the final valid descriptor is always programmed to generate an interrupt.

There are two general methods for managing a descriptor queue using interrupts:

- Interrupt on every descriptor
- Interrupt minimally - only on the last descriptor

Descriptor Queue Using Interrupts on Every Descriptor

In this system, the DMA Manager software synchronizes with the DMA unit by enabling an interrupt on every descriptor. This method should only be used if system design can guarantee that each interrupt event will be serviced separately (no interrupt overrun).

To maintain synchronization of the descriptor queue, the non-interrupt software maintains a count of descriptors added to the queue, while the interrupt handler maintains a count of completed descriptors removed from the queue. The counts are equal only when the DMA channel is paused after having processed all the descriptors.

When each new work request is received, the DMA Manager software initializes a new descriptor, taking care to write a `DMA_CONFIG` value with a `FLOW` value of 0. Next, the software compares the descriptor counts to determine if the DMA channel is running or not. If the DMA channel is paused (counts equal), the software increments its count and then starts the DMA unit by writing the new descriptor's `DMA_CONFIG` value to the DMA channel's `DMA_CONFIG` register.

If the counts are unequal, the software instead modifies the next-to-last descriptor's DMA_CONFIG value so that its upper half (FLOW and NDSIZE) now describes the newly queued descriptor. This operation does not disrupt the DMA channel, provided the rest of the descriptor data structure is initialized in advance. It is necessary, however, to synchronize the software to the DMA to correctly determine whether the new or the old DMA_CONFIG value was read by the DMA channel.

This synchronization operation should be performed in the interrupt handler. First, upon interrupt, the handler should read the channel's IRQ_STATUS register. If the DMA_RUN status bit is set, then the channel has moved on to processing another descriptor, and the interrupt handler may increment its count and exit. If the DMA_RUN status bit is not set, however, then the channel has paused, either because there are no more descriptors to process, or because the last descriptor was queued too late (that is, the modification of the next-to-last descriptor's DMA_CONFIG element occurred after that element was read into the DMA unit.) In this case, the interrupt handler should write the DMA_CONFIG value appropriate for the last descriptor to the DMA channel's DMA_CONFIG register, increment the completed descriptor count, and exit.

Again, this system can fail if the system's interrupt latencies are large enough to cause any of the channel's DMA interrupts to be dropped. An interrupt handler capable of safely synchronizing multiple descriptors' interrupts would need to be complex, performing several MMR accesses to ensure robust operation. In such a system environment, a minimal interrupt synchronization method is preferred.

Descriptor Queue Using Minimal Interrupts

In this system, only one DMA interrupt event is possible in the queue at any time. The DMA interrupt handler for this system can also be extremely short. Here, the descriptor queue is organized into an "active" and a "waiting" portion, where interrupts are enabled only on the last descriptor in each portion.

When each new DMA request is processed, the software's non-interrupt code fills in a new descriptor's contents and adds it to the waiting portion of the queue. The descriptor's DMA_CONFIG word should have a FLOW value of zero. If more than one request is received before the DMA queue completion interrupt occurs, the non-interrupt code should queue later descriptors, forming a waiting portion of the queue that is disconnected from the active portion of the queue being processed by the DMA unit. In other words, all but the last active descriptors contain FLOW values >= 4 and have no interrupt enable set, while the last active descriptor contains a FLOW of 0 and an interrupt enable bit DI_EN set to 1. Also, all but the last waiting descriptors contain FLOW values >= 4 and no interrupt enables set, while the last waiting descriptor contains a FLOW of 0 and an interrupt enable bit set to 1. This ensures that the DMA unit can automatically process the whole active queue and then issue one interrupt. Also, this arrangement makes it easy to start the waiting queue within the interrupt handler by a single DMA_CONFIG register write.

After queuing a new waiting descriptor, the non-interrupt software should leave a message for its interrupt handler in a memory mailbox location containing the desired DMA_CONFIG value to use to start the first waiting descriptor in the waiting queue (or 0 to indicate no descriptors are waiting.)

It is critical that the software not modify the contents of the active descriptor queue directly, once its processing by the DMA unit has been started, unless careful synchronization measures are taken. In the most straightforward implementation of a descriptor queue, the DMA Manager software would never modify descriptors on the active queue; instead, the DMA Manager waits until the DMA queue completion interrupt indicates the processing of the entire active queue is complete.

When a DMA queue completion interrupt is received, the interrupt handler reads the mailbox from the non-interrupt software and writes the value in it to the DMA channel's DMA_CONFIG register. This single register write restarts the queue, effectively transforming the waiting queue to an

active queue. The interrupt handler should then pass a message back to the non-interrupt software indicating the location of the last descriptor accepted into the active queue. If, on the other hand, the interrupt handler reads its mailbox and finds a DMA_CONFIG value of zero, indicating there is no more work to perform, then it should pass an appropriate message (for example, zero) back to the non-interrupt software indicating that the queue has stopped. This simple handler should be able to be coded in a very small number of instructions.

The non-interrupt software which accepts new DMA work requests needs to synchronize the activation of new work with the interrupt handler. If the queue has stopped (that is, if the mailbox from the interrupt software is zero), the non-interrupt software is responsible for starting the queue (writing the first descriptor's DMA_CONFIG value to the channel's DMA_CONFIG register). If the queue is not stopped, however, the non-interrupt software must not write the DMA_CONFIG register (which would cause a DMA error), but instead it should queue the descriptor onto the waiting queue and update its mailbox directed to the interrupt handler.

DMA Errors (Aborts)

The DMA controller flags conditions that cause the DMA process to end abnormally (that is, abort). This functionality is provided as a tool for system development and debug, as a way to detect DMA-related programming errors. DMA errors (aborts) are detected by the DMA channel module in the cases listed below. When a DMA error occurs, the channel is immediately stopped (DMA_RUN goes to 0) and any prefetched data is discarded. In addition, a DMA_ERROR interrupt is asserted.

There is only one DMA_ERROR interrupt for the whole DMA controller, which is asserted whenever any of the channels has detected an error condition.

The DMA_ERROR interrupt handler must do these things for each channel:

- Read each channel's IRQ_STATUS register to look for a channel with the DMA_ERR bit set (bit 1).
- Disable the channel with DMA_ERR bit set (write 0 to DMA_CONFIG register bit 0).
- Clear the DMA_ERR bit (write IRQ_STATUS with bit 1 = 1).
- Clear the problem with that channel (for example, fix register values).
- Followed by SSYNC.

The following error conditions are detected by the DMA hardware and result in a DMA Abort interrupt.

- The Configuration register contains invalid values:

  - Incorrect WDSIZE value (WDSIZE = b#11)

  - Bit 15 not set to 0

  - Incorrect FLOW value (FLOW = 2, 3, or 5)

  - NDSIZE value does not agree with FLOW. See Table 9-8.

Table 9-8. Legal NDSIZE Values

FLOW	NDSIZE	Note
0	0	
1	0	
4	0 < NDSIZE <= 7	Descriptor array, no descriptor pointer fetched
6	0 < NDSIZE <= 8	Descriptor list, small descriptor pointer fetched
7	0 < NDSIZE <= 9	Descriptor list, large descriptor pointer fetched

- A disallowed register write occurred while the channel was running. Only the DMA_CONFIG and IRQ_STATUS registers can be written when DMA_RUN = 1.
- An address alignment error occurred during any memory access. For example, DMA_CONFIG register WDSIZE = 1 (16 bit) but the least significant bit (LSB) of the address is not equal to 0, or WDSIZE = 2 (32 bit) but the two LSBs of the address are not equal to 00.
- A memory space transition was attempted (internal-to-external or vice versa).
- A memory access error occurred. Either an access attempt was made to an internal address not populated or defined as cache, or an external access caused an error (signaled by the external memory interface).

Some prohibited situations are not detected by the DMA hardware. No DMA abort is signaled for these situations:

- DMA_CONFIG Direction bit (WNR) does not agree with the direction of the mapped peripheral.
- DMA_CONFIG Direction bit does not agree with the direction of the MDMA channel.

- DMA_CONFIG Word Size (WDSIZE) is not supported by the mapped peripheral.
- DMA_CONFIG Word Size in source and destination of the MDMA stream are not equal.
- Descriptor chain indicates data buffers that are not in the same internal/external memory space.
- In 2D DMA, X_COUNT = 1.

10 SPI COMPATIBLE PORT CONTROLLERS

The processor has a Serial Peripheral Interface (SPI) port that provides an I/O interface to a wide variety of SPI compatible peripheral devices.

With a range of configurable options, the SPI port provides a glueless hardware interface with other SPI compatible devices. SPI is a four-wire interface consisting of two data pins, a device select pin, and a clock pin. SPI is a full-duplex synchronous serial interface, supporting master modes, slave modes, and multimaster environments. The SPI compatible peripheral implementation also supports programmable baud rate and clock phase/polarities. The SPI features the use of open drain drivers to support the multimaster scenario and to avoid data contention.

Typical SPI compatible peripheral devices that can be used to interface to the SPI compatible interface include:

- Other CPUs or microcontrollers
- Codecs
- A/D converters
- D/A converters
- Sample rate converters
- SP/DIF or AES/EBU digital audio transmitters and receivers
- LCD displays

- Shift registers
- FPGAs with SPI emulation

The SPI is an industry-standard synchronous serial link that supports communication with multiple SPI compatible devices. The SPI peripheral is a synchronous, four-wire interface consisting of two data pins (MOSI and MISO), one device select pin ($\overline{\text{SPISS}}$), and a gated clock pin (SCK). With the two data pins, it allows for full-duplex operation to other SPI compatible devices. The SPI also includes programmable baud rates, clock phase, and clock polarity.

The SPI can operate in a multimaster environment by interfacing with several other devices, acting as either a master device or a slave device. In a multimaster environment, the SPI interface uses open drain outputs to avoid data bus contention.

Figure 10-1 provides a block diagram of the SPI. The interface is essentially a shift register that serially transmits and receives data bits, one bit at a time at the SCK rate, to and from other SPI devices. SPI data is transmitted and received at the same time through the use of a shift register. When an SPI transfer occurs, data is simultaneously transmitted (shifted serially out of the shift register) as new data is received (shifted serially into the other end of the same shift register). The SCK synchronizes the shifting and sampling of the data on the two serial data pins.

During SPI data transfers, one SPI device acts as the SPI link master, where it controls the data flow by generating the SPI serial clock and asserting the SPI device select signal ($\overline{\text{SPISS}}$). The other SPI device acts as the slave and accepts new data from the master into its shift register, while it transmits requested data out of the shift register through its SPI transmit data pin. Multiple processors can take turns being the master device, as can other microcontrollers or microprocessors. One master device can also simultaneously shift data into multiple slaves (known as Broadcast mode). However, only one slave may drive its output to write data back to the master at any given time. This must be enforced in Broadcast mode,

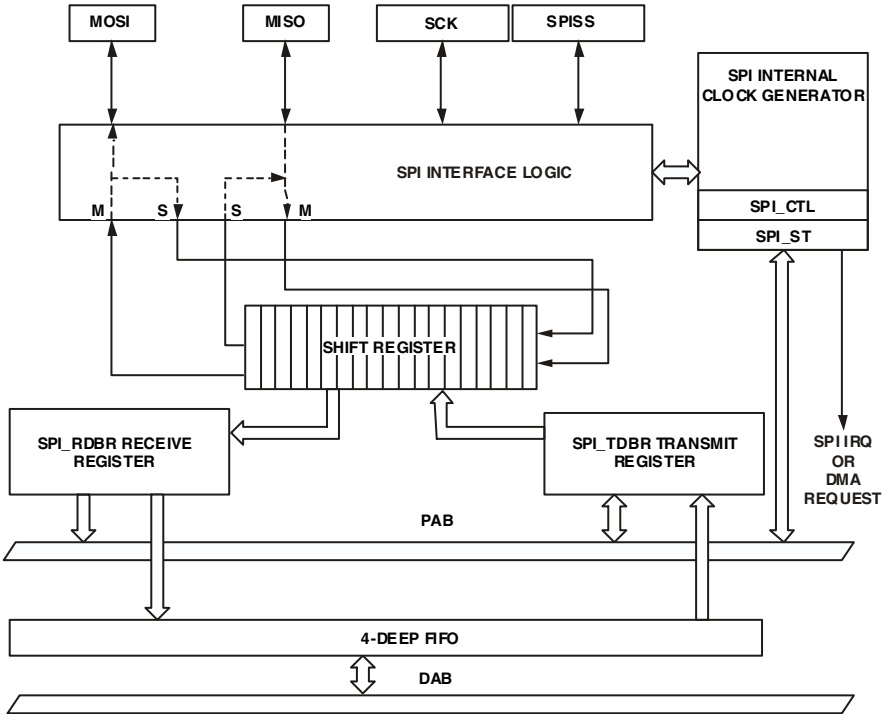

Figure 10-1. SPI Block Diagram

where several slaves can be selected to receive data from the master, but only one slave at a time can be enabled to send data back to the master.

In a multimaster or multidevice environment where multiple processors are connected via their SPI ports, all MOSI pins are connected together, all MISO pins are connected together, and all SCK pins are connected together.

For a multislave environment, the processor can make use of seven programmable flags, PF1–PF7, that are dedicated SPI slave select signals for the SPI slave devices.

At reset, the SPI is disabled and configured as a slave.

Interface Signals

The following section discusses the SPI signals.

Serial Peripheral Interface Clock Signal (SCK)

The SCK signal is the SPI clock signal. This control signal is driven by the master and controls the rate at which data is transferred. The master may transmit data at a variety of baud rates. The SCK signal cycles once for each bit transmitted. It is an output signal if the device is configured as a master, and an input signal if the device is configured as a slave.

The SCK is a gated clock that is active during data transfers only for the length of the transferred word. The number of active clock edges is equal to the number of bits driven on the data lines. Slave devices ignore the serial clock if the Serial Peripheral Slave Select Input ($\overline{\text{SPISS}}$) is driven inactive (high).

The SCK is used to shift out and shift in the data driven on the MISO and MOSI lines. Clock polarity and clock phase relative to data are programmable in the SPI Control register (SPI_CTL) and define the transfer format. (See "SPI Transfer Protocols" on page 10-20.)

Serial Peripheral Interface Slave Select Input Signal

The $\overline{\text{SPISS}}$ signal is the SPI Serial Peripheral Slave Select Input signal. This is an active-low signal used to enable a processor when it is configured as a slave device. This input-only pin behaves like a chip select and is provided by the master device for the slave devices. For a master device, it can act as an error signal input in case of the multimaster environment. In multimaster mode, if the $\overline{\text{SPISS}}$ input signal of a master is asserted (driven low), and the PSSE bit in the SPI_CTL register is enabled, an error

has occurred. This means that another device is also trying to be the master device.

The $\overline{\text{SPISS}}$ signal is multiplexed with the PF0/TMR0 pin.

Master Out Slave In (MOSI)

The MOSI pin is the Master Out Slave In pin, one of the bidirectional I/O data pins. If the processor is configured as a master, the MOSI pin becomes a data transmit (output) pin, transmitting output data. If the processor is configured as a slave, the MOSI pin becomes a data receive (input) pin, receiving input data. In an SPI interconnection, the data is shifted out from the MOSI output pin of the master and shifted into the MOSI input(s) of the slave(s).

Master In Slave Out (MISO)

The MISO pin is the Master In Slave Out pin, one of the bidirectional I/O data pins. If the processor is configured as a master, the MISO pin becomes a data receive (input) pin, receiving input data. If the processor is configured as a slave, the MISO pin becomes a data transmit (output) pin, transmitting output data. In an SPI interconnection, the data is shifted out from the MISO output pin of the slave and shifted into the MISO input pin of the master.

Only one slave is allowed to transmit data at any given time.

The SPI configuration example in Figure 10-2 illustrates how the processor can be used as the slave SPI device. The 8-bit host microcontroller is the SPI master.

The processor can be booted via its SPI interface to allow user application code and data to be downloaded before runtime.

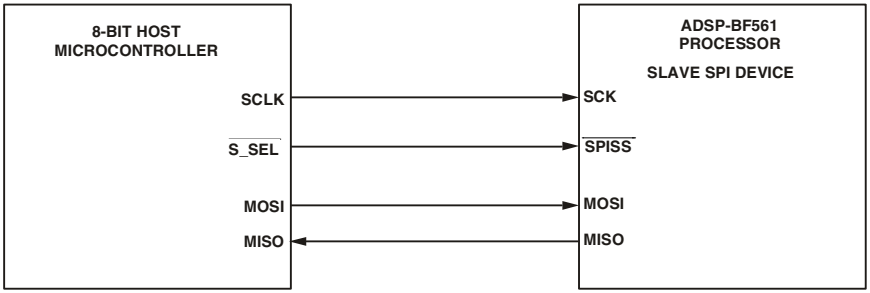

Figure 10-2. ADSP-BF561 Processor as Slave SPI Device

Interrupt Output

The SPI has two interrupt output signals: a data interrupt and an error interrupt.

The behavior of the SPI data interrupt signal depends on the Transfer Initiation mode bit field (TIMOD) in the SPI Control register. In DMA mode (TIMOD = 1X), the data interrupt acts as a DMA request and is generated when the DMA FIFO is ready to be written to (TIMOD = 11) or read from (TIMOD = 10). In non-DMA mode (TIMOD = 0X), a data interrupt is generated when the SPI_TDBR is ready to be written to (TIMOD = 01) or when the SPI_RDBR is ready to be read from (TIMOD = 00).

An SPI Error interrupt is generated in a master when a Mode Fault Error occurs, in both DMA and non-DMA modes. An error interrupt can also be generated in DMA mode when there is an underflow (TXE when TIMOD = 11) or an overflow (RBSY when TIMOD = 10) error condition. In non-DMA mode, the underflow and overflow conditions set the TXE and RBSY bits in the SPI_STAT register, respectively, but do not generate an error interrupt.

For more information about this interrupt output, see the discussion of the TIMOD bits in “SPI Control Register (SPI_CTL)” on page 10-8.

SPI Registers

The SPI peripheral includes a number of user-accessible registers. Some of these registers are also accessible through the DMA bus. Four registers contain control and status information: SPI_BAUD, SPI_CTL, SPI_FLG, and SPI_STAT. Two registers are used for buffering receive and transmit data: SPI_RDBR and SPI_TDBR. For information about DMA-related registers, see Chapter 9, "Direct Memory Access." The shift register, SFDR, is internal to the SPI module and is not directly accessible.

See "Error Signals and Flags" on page 10-27 for more information about how the bits in these registers are used to signal errors and other conditions. See "Register Functions" on page 10-19 for more information about SPI register and bit functions.

SPI Baud Rate Register (SPI_BAUD)

The SPI Baud Rate register (SPI_BAUD) is used to set the bit transfer rate for a master device (see Figure 10-3). When configured as a slave, the value written to this register is ignored. The SPI clock frequency is determined by this formula:

$$SCK\ Frequency = (Peripheral\ clock\ frequency\ SCLK)/(2 \times SPI_BAUD)$$

Writing a value of 0 or 1 to the register disables the SPI clock. Therefore, the maximum SPI clock rate is one-fourth the system clock rate.

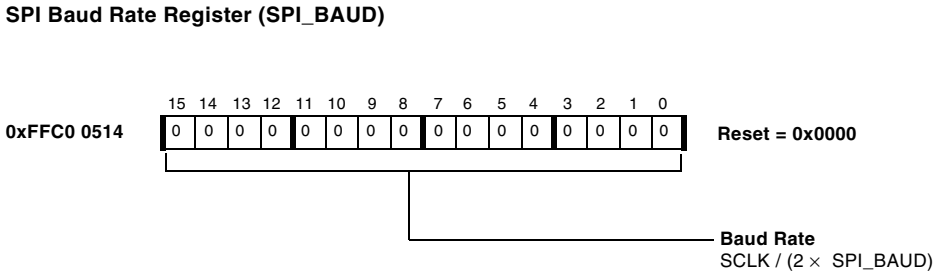

Figure 10-3. SPI Baud Rate Register

Table 10-1 lists several possible baud rate values for SPI_BAUD.

Table 10-1. SPI Master Baud Rate Example

SPI_BAUD Decimal Value	SPI Clock (SCK) Divide Factor	Baud Rate for SCLK at 100 MHz
0	N/A	N/A
1	N/A	N/A
2	4	25 MHz
3	6	16.7 MHz
4	8	12.5 MHz
65,535 (0xFFFF)	131,070	763 Hz

SPI Control Register (SPI_CTL)

The SPI_CTL register is used to configure and enable the SPI system. This register is used to enable the SPI interface, select the device as a master or slave, and determine the data transfer format and word size.

The term “word” refers to a single data transfer of either 8 bits or 16 bits, depending on the word length (SIZE) bit in SPI_CTL. There are two special bits which can also be modified by the hardware: SPE and MSTR.

The TIMOD field is used to specify the action that initiates transfers to/from the receive/transmit buffers. When set to 00, a SPI port transaction is begun when the receive buffer is read. Data from the first read will need to be discarded since the read is needed to initiate the first SPI port transaction. When set to 01, the transaction is initiated when the transmit buffer is written. A value of 10 selects DMA Receive mode and the first transaction is initiated by enabling the SPI for DMA Receive mode. Subsequent individual transactions are initiated by a DMA read of the SPI_RDBR. A value of 11 selects DMA Transmit mode and the transaction is initiated by a DMA write of the SPI_TDBR.

The PSSE bit is used to enable the $\overline{\text{SPISS}}$ input for master. When not used, $\overline{\text{SPISS}}$ can be disabled, freeing up a chip pin as general-purpose I/O.

The EMISO bit enables the MISO pin as an output. This is needed in an environment where the master wishes to transmit to various slaves at one time (broadcast). Only one slave is allowed to transmit data back to the master. Except for the slave from whom the master wishes to receive, all other slaves should have this bit cleared.

The SPE and MSTR bits can be modified by hardware when the MODF bit of the Status register is set. See “Mode Fault Error” on page 10-27.

Figure 10-4 provides the bit descriptions for SPI_CTL.

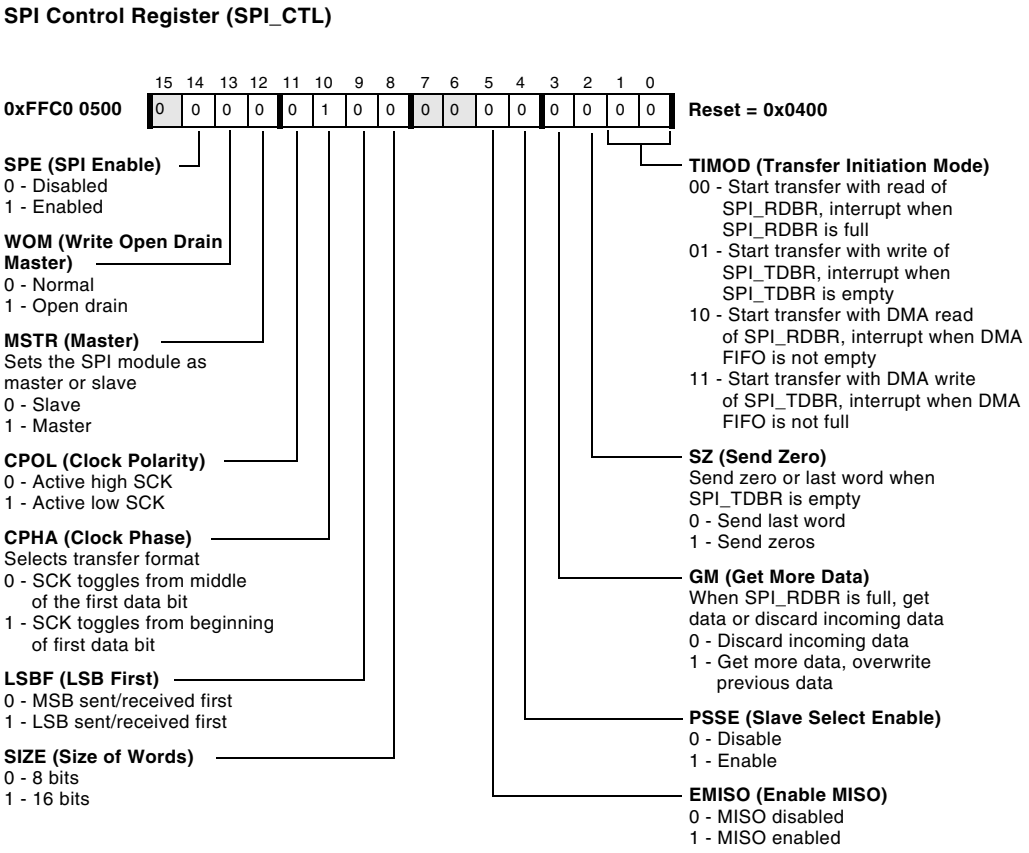

Figure 10-4. SPI Control Register

SPI Flag Register (SPI_FLG)

If the SPI is enabled as a master, the SPI uses SPI_FLG to enable up to seven general-purpose programmable flag pins to be used as individual slave select lines. In Slave mode, the SPI_FLG bits have no effect, and each SPI uses the $\overline{\text{SPISS}}$ input as a slave select. Figure 10-5 shows the SPI_FLG register diagram.

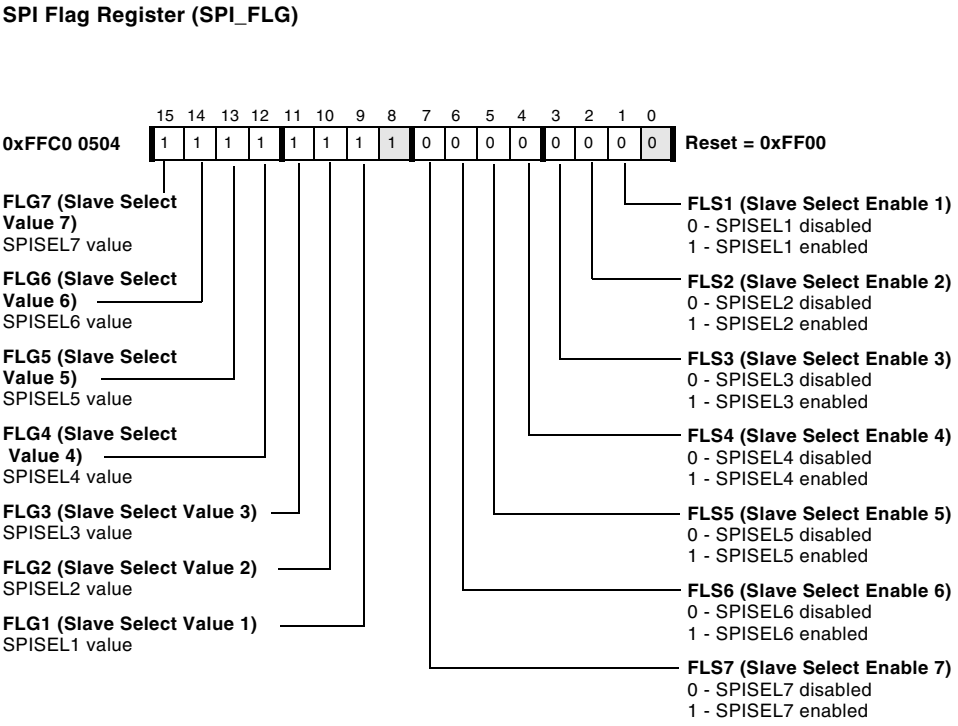

Figure 10-5. SPI Flag Register

The SPI_FLG register consists of two sets of bits that function as follows.

- Slave Select Enable (FLSx) bits

  Each FLSx bit corresponds to a Programmable Flag (PFx) pin. When a FLSx bit is set, the corresponding PFx pin is driven as a slave select. For example, if FLS1 is set in SPI_FLG, PF1 is driven as a slave select (SPISEL1). Table 10-2 shows the association of the FLSx bits and the corresponding PFx pins.

  If the FLSx bit is not set, the general-purpose programmable flag registers (FIO_DIR and others) configure and control the corresponding PFx pin.

- Slave Select Value (FLGx) bits

  When a PFx pin is configured as a slave select output, the FLGx bits can determine the value driven onto the output. If the CPHA bit in SPI_CTL is set, the output value is set by software control of the FLGx bits. The SPI protocol permits the slave select line to either remain asserted (low) or be deasserted between transferred words. The user must set or clear the appropriate FLGx bits. For example, to drive PF3 as a slave select, FLS3 in SPI_FLG must be set. Clearing FLG3 in SPI_FLG drives PF3 low; setting FLG3 drives PF3 high. The PF3 pin can be cycled high and low between transfers by setting and clearing FLG3. Otherwise, PF3 remains active (low) between transfers.

  If CPHA = 0, the SPI hardware sets the output value and the FLGx bits are ignored. The SPI protocol requires that the slave select be deasserted between transferred words. In this case, the SPI hardware controls the pins. For example, to use PF3 as a slave select pin, it is only necessary to set the FLS3 bit in SPI_FLG. It is not necessary to write to the FLG3 bit, because the SPI hardware automatically drives the PF3 pin.

Table 10-2. SPI_FLG Bit Mapping to PFx Pins

Bit	Name	Function	PFx Pin	Default
0		Reserved		0
1	FLS1	SPISEL1 Enable	PF1	0
2	FLS2	SPISEL2 Enable	PF2	0
3	FLS3	SPISEL3 Enable	PF3	0
4	FLS4	SPISEL4 Enable	PF4	0
5	FLS5	SPISEL5 Enable	PF5	0
6	FLS6	SPISEL6 Enable	PF6	0
7	FLS7	SPISEL7 Enable	PF7	0
8		Reserved		1
9	FLG1	SPISEL1 Value	PF1	1
10	FLG2	SPISEL2 Value	PF2	1
11	FLG3	SPISEL3 Value	PF3	1
12	FLG4	SPISEL4 Value	PF4	1
13	FLG5	SPISEL5 Value	PF5	1
14	FLG6	SPISEL6 Value	PF6	1
15	FLG7	SPISEL7 Value	PF7	1

Slave Select Inputs

If the SPI is in Slave mode, $\overline{\texttt{SPISS}}$ acts as the slave select input. When enabled as a master, $\overline{\texttt{SPISS}}$ can serve as an error detection input for the SPI in a multimaster environment. The `PSSE` bit in `SPI_CTL` enables this feature. When `PSSE = 1`, the $\overline{\texttt{SPISS}}$ input is the master mode error input. Otherwise, $\overline{\texttt{SPISS}}$ is ignored.

Use of FLS Bits in SPI_FLG for Multiple Slave SPI Systems

The FLSx bits in the SPI_FLG register are used in a multiple slave SPI environment. For example, if there are eight SPI devices in the system including a processor master, the master processor can support the SPI mode transactions across the other seven devices. This configuration requires only one master processor in this multislave environment. For example, assume that the SPI is the master. The seven flag pins (PF1-PF7) on the processor can be connected to each of the slave SPI device's $\overline{\text{SPISS}}$ pins. In this configuration, the FLSx bits in SPI_FLG can be used in three cases.

In cases 1 and 2, the processor is the master and the seven microcontrollers/peripherals with SPI interfaces are slaves. The processor can:

1. Transmit to all seven SPI devices at the same time in a broadcast mode. Here, all FLSx bits are set.
2. Receive and transmit from one SPI device by enabling only one slave SPI device at a time.

   In case 3, all eight devices connected via SPI ports can be other processors.
3. If all the slaves are also processors, then the requester can receive data from only one processor (enabled by clearing the EMISO bit in the six other slave processors) at a time and transmit broadcast data to all seven at the same time. This EMISO feature may be available in some other microcontrollers. Therefore, it is possible to use the EMISO feature with any other SPI device that includes this functionality.

Figure 10-6 shows one processor as a master with three processors (or other SPI compatible devices) as slaves.

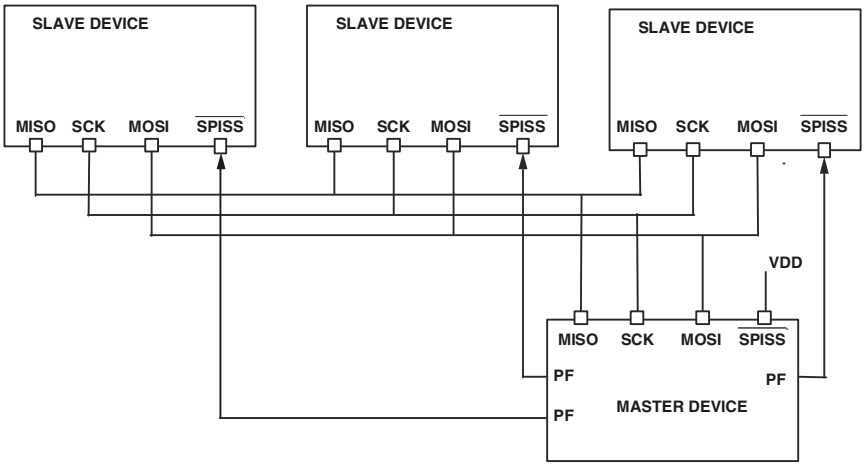

Figure 10-6. Single-Master, Multiple-Slave Configuration

SPI Status Register (SPI_STAT)

The SPI Status register (SPI_STAT) is used to detect when an SPI transfer is complete or if transmission/reception errors occur (Figure 10-7). The SPI_STAT register can be read at any time.

Some of the bits in SPI_STAT are read-only and other bits are sticky. Bits that provide information only about the SPI are read-only. These bits are set and cleared by the hardware. Sticky bits are set when an error condition occurs. These bits are set by hardware and must be cleared by software. To clear a sticky bit, the user must write a 1 to the desired bit position of SPI_STAT. For example, if the TXE bit is set, the user must write a 1 to bit 2 of SPI_STAT to clear the TXE error condition. This allows the user to read SPI_STAT without changing its value.

Sticky bits are cleared on a reset, but are not cleared on an SPI disable.

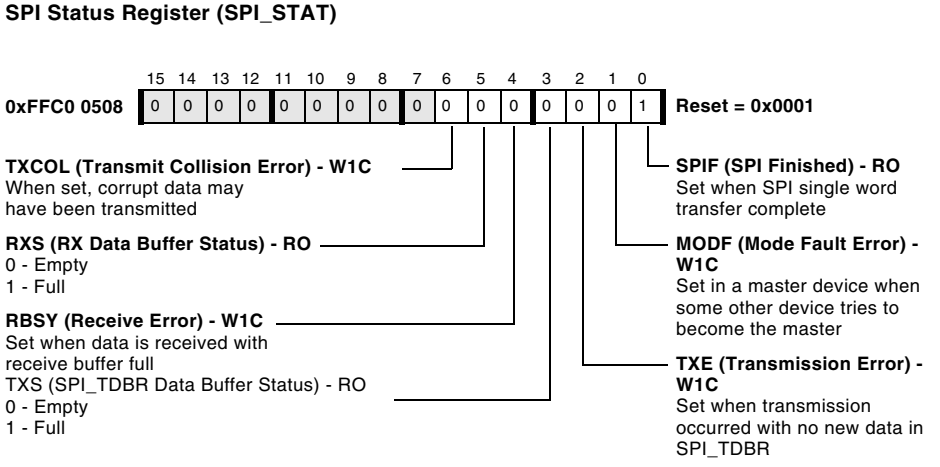

Figure 10-7. SPI Status Register

The transmit buffer becomes full after it is written to. It becomes empty when a transfer begins and the transmit value is loaded into the shift register. The receive buffer becomes full at the end of a transfer when the shift register value is loaded into the receive buffer. It becomes empty when the receive buffer is read.

The SPIF bit is set when the SPI port is disabled.

Upon entering DMA mode, the transmit buffer and the receive buffer become empty. That is, the TXS bit and the RXS bit are initially cleared upon entering DMA mode.

When using DMA for SPI transmit, the DMA_DONE interrupt signifies that the DMA FIFO is empty. However, at this point there may still be data in the SPI DMA FIFO waiting to be transmitted. Therefore, software needs to poll TXS in the SPI_STAT register until it goes low for 2 successive reads, at which point the SPI DMA FIFO will be empty. When the SPIF bit subsequently goes low, the last word has been transferred and the SPI can be disabled or enabled for another mode.

SPI Transmit Data Buffer Register (SPI_TDBR)

The SPI Transmit Data Buffer register (SPI_TDBR) is a 16-bit read-write register (Figure 10-8). Data is loaded into this register before being transmitted. Just prior to the beginning of a data transfer, the data in SPI_TDBR is loaded into a shift register. A read of SPI_TDBR can occur at any time and does not interfere with or initiate SPI transfers.

When the DMA is enabled for transmit operation, the DMA engine loads data into this register for transmission just prior to the beginning of a data transfer. A write to SPI_TDBR should not occur in this mode because this data will overwrite the DMA data to be transmitted.

When the DMA is enabled for receive operation, the contents of SPI_TDBR are repeatedly transmitted. A write to SPI_TDBR is permitted in this mode, and this data is transmitted.

If the Send Zeros control bit (SZ in the SPI_CTL register) is set, SPI_TDBR may be reset to 0 under certain circumstances.

If multiple writes to SPI_TDBR occur while a transfer is already in progress, only the last data written is transmitted. None of the intermediate values written to SPI_TDBR are transmitted. Multiple writes to SPI_TDBR are possible, but not recommended.

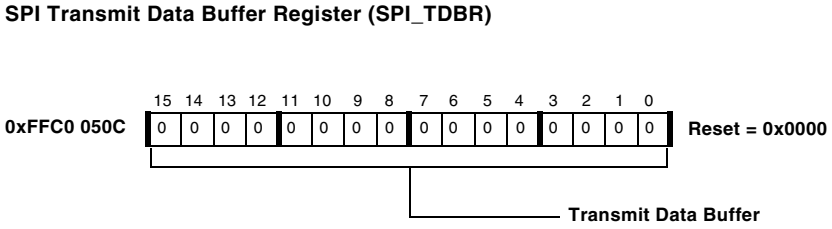

Figure 10-8. SPI Transmit Data Buffer Register

SPI Receive Data Buffer Register (SPI_RDBR)

The SPI Receive Data Buffer register (SPI_RDBR) is a 16-bit read-only register (Figure 10-9). At the end of a data transfer, the data in the shift register is loaded into SPI_RDBR. During a DMA receive operation, the data in SPI_RDBR is automatically read by the DMA. When SPI_RDBR is read via software, the RXS bit is cleared and an SPI transfer may be initiated (if TIMOD = 00).

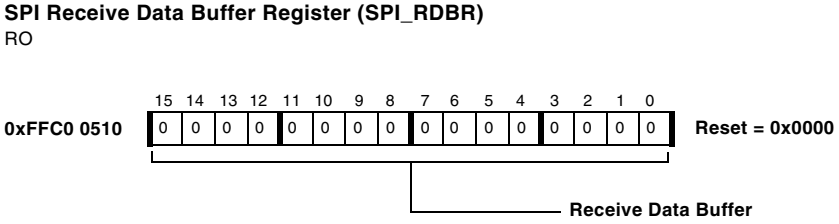

Figure 10-9. SPI Receive Data Buffer Register

SPI RDBR Shadow Register (SPI_SHADOW)

A shadow register for the receive data buffer, SPI_RDBR, has been provided for use in debugging software (Figure 10-10). This register is at a different address than SPI_RDBR, but its contents are identical to that of SPI_RDBR. When a software read of SPI_RDBR occurs, the RXS bit in SPI_STAT is cleared and an SPI transfer may be initiated (if TIMOD = 00 in SPI_CTL). No such hardware action occurs when the SPI_SHADOW register is read. The SPI_SHADOW register is read-only.

SPI RDBR Shadow Register (SPI_SHADOW)
RO

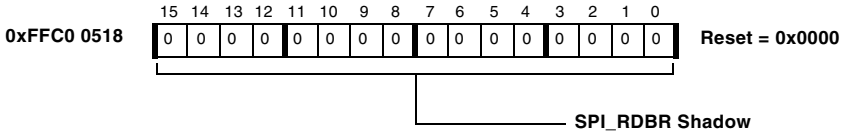

Figure 10-10. SPI RDBR Shadow Register

Register Functions

Table 10-3 summarizes the functions of the SPI registers.

Table 10-3. SPI Register Mapping

Register Name	Function	Notes
SPI_CTL	SPI port control	SPE and MSTR bits can also be modified by hardware (when MODF is set).
SPI_FLG	SPI port flag	Bits 0 and 8 are reserved.
SPI_STAT	SPI port status	SPIF bit can be set by clearing SPE in SPI_CTL.
SPI_TDBR	SPI port transmit data buffer	Register contents can also be modified by hardware (by DMA and/or when SZ = 1 in SPI_CTL).
SPI_RDBR	SPI port receive data buffer	When register is read, hardware events may be triggered.
SPI_BAUD	SPI port baud control	Value of 0 or 1 disables the SPI clock.
SPI_SHADOW	SPI port data	Register has the same contents as SPI_RDBR, but no hardware events are triggered when read.

SPI Transfer Protocols

The SPI supports four different combinations of serial clock phase and polarity (SPI modes 0-3). These combinations are selected using the CPOL and CPHA bits in SPI_CTL, as shown in Figure 10-11.

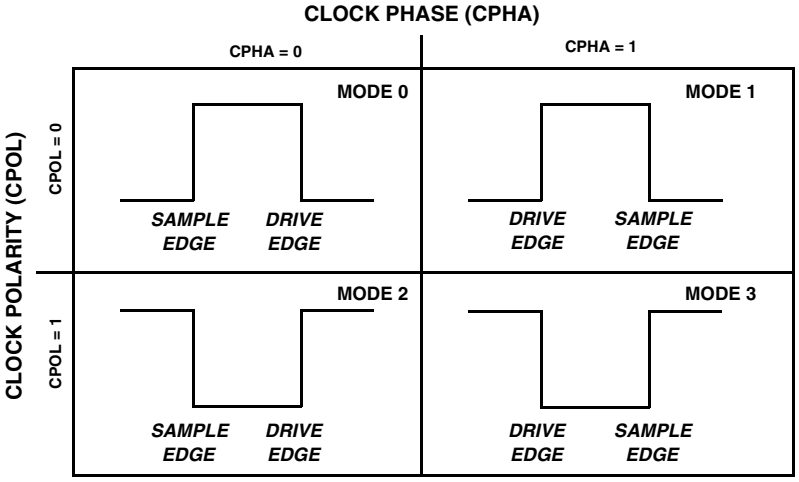

Figure 10-11. SPI Modes of Operation

The figures "SPI Transfer Protocol for CPHA = 0" on page 10-21 and "SPI Transfer Protocol for CPHA = 1" on page 10-22 demonstrate the two basic transfer formats as defined by the CPHA bit. Two waveforms are shown for SCK—one for CPOL = 0 and the other for CPOL = 1. The diagrams may be interpreted as master or slave timing diagrams since the SCK, MISO, and MOSI pins are directly connected between the master and the slave. The MISO signal is the output from the slave (slave transmission), and the MOSI signal is the output from the master (master transmission). The SCK signal is generated by the master, and the $\overline{\text{SPISS}}$ signal is the slave device select input to the slave from the master. The diagrams represent an 8-bit transfer (SIZE = 0) with the Most Significant Bit (MSB) first (LSBF = 0). Any combination of the SIZE and LSBF bits of SPI_CTL is allowed. For

example, a 16-bit transfer with the Least Significant Bit (LSB) first is another possible configuration.

The clock polarity and the clock phase should be identical for the master device and the slave device involved in the communication link. The transfer format from the master may be changed between transfers to adjust to various requirements of a slave device.

When CPHA = 0, the slave select line, $\overline{\text{SPISS}}$, must be inactive (high) between each serial transfer. This is controlled automatically by the SPI hardware logic. When CPHA = 1, $\overline{\text{SPISS}}$ may either remain active (low) between successive transfers or be inactive (high). This must be controlled by the software via manipulation of SPI_FLG.

Figure 10-12 shows the SPI transfer protocol for CPHA = 0. Note SCK starts toggling in the middle of the data transfer, SIZE = 0, and LSBF = 0.

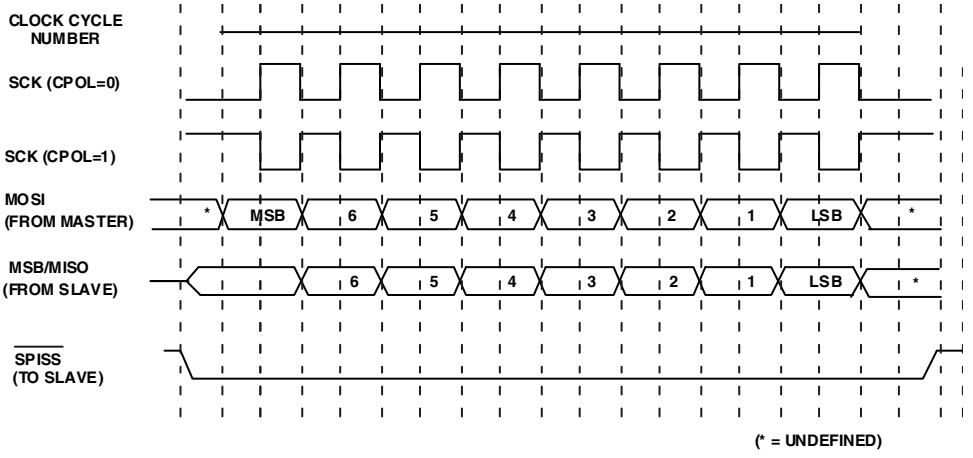

Figure 10-12. SPI Transfer Protocol for CPHA = 0

Figure 10-13 shows the SPI transfer protocol for CPHA = 1. Note SCK starts toggling at the beginning of the data transfer, SIZE = 0, and LSBF = 0.

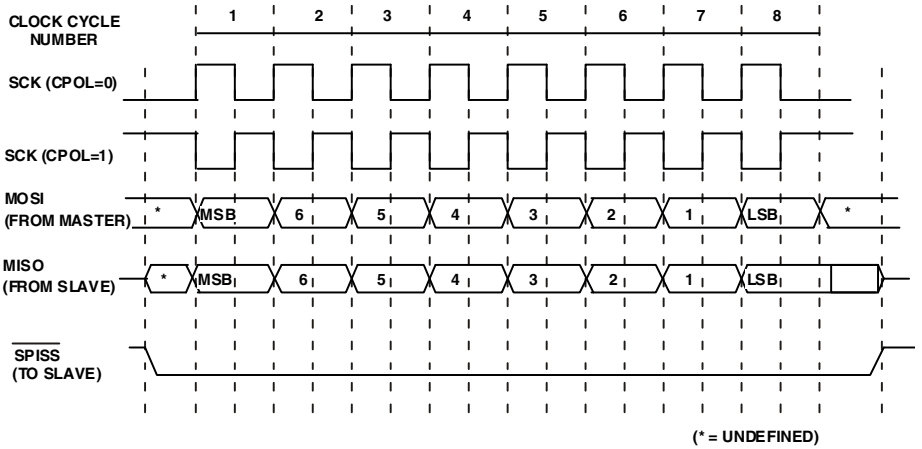

Figure 10-13. SPI Transfer Protocol for CPHA = 1

SPI General Operation

The SPI can be used in a single master as well as multimaster environment. The MOSI, MISO, and the SCK signals are all tied together in both configurations. SPI transmission and reception are always enabled simultaneously, unless the Broadcast mode has been selected. In Broadcast mode, several slaves can be enabled to receive, but only one of the slaves must be in Transmit mode driving the MISO line. If the transmit or receive is not needed, it can simply be ignored. This section describes the clock signals, SPI operation as a master and as a slave, and error generation.

Precautions must be taken to avoid data corruption when changing the SPI module configuration. The configuration must not be changed during a data transfer. The clock polarity should only be changed when no slaves are selected. An exception to this is when an SPI communication link consists of a single master and a single slave, CPHA = 1, and the slave select

input of the slave is always tied low. In this case, the slave is always selected and data corruption can be avoided by enabling the slave only after both the master and slave devices are configured.

In a multimaster or multislave SPI system, the data output pins (MOSI and MISO) can be configured to behave as open drain outputs, which prevents contention and possible damage to pin drivers. An external pull-up resistor is required on both the MOSI and MISO pins when this option is selected.

The WOM bit controls this option. When WOM is set and the SPI is configured as a master, the MOSI pin is three-stated when the data driven out on MOSI is a logic high. The MOSI pin is not three-stated when the driven data is a logic low. Similarly, when WOM is set and the SPI is configured as a slave, the MISO pin is three-stated if the data driven out on MISO is a logic high.

Clock Signals

The SCK signal is a gated clock that is only active during data transfers for the duration of the transferred word. The number of active edges is equal to the number of bits driven on the data lines. The clock rate can be as high as one-fourth of the SCLK rate. For master devices, the clock rate is determined by the 16-bit value of SPI_BAUD. For slave devices, the value in SPI_BAUD is ignored. When the SPI device is a master, SCK is an output signal. When the SPI is a slave, SCK is an input signal. Slave devices ignore the SPI clock if the slave select input is driven inactive (high).

The SCK signal is used to shift out and shift in the data driven onto the MISO and MOSI lines. The data is always shifted out on one edge of the clock and sampled on the opposite edge of the clock. Clock polarity and clock phase relative to data are programmable into SPI_CTL and define the transfer format. (See Figure 10-11.)

Master Mode Operation

When the SPI is configured as a master (and DMA mode is not selected), the interface operates in the following manner.

1. The core writes to SPI_FLG, setting one or more of the SPI Flag Select bits (FLSx). This ensures that the desired slaves are properly deselected while the master is configured.
2. The core writes to the SPI_BAUD and SPI_CTL registers, enabling the device as a master and configuring the SPI system by specifying the appropriate word length, transfer format, baud rate, and other necessary information.
3. If CPHA = 1, the core activates the desired slaves by clearing one or more of the SPI flag bits (FLGx) of SPI_FLG.
4. The TIMOD bits in SPI_CTL determine the SPI Transfer Initiate mode. The transfer on the SPI link begins upon either a data write by the core to the transmit data buffer (SPI_TDBR) or a data read of the receive data buffer (SPI_RDBR).
5. The SPI then generates the programmed clock pulses on SCK and simultaneously shifts data out of MOSI and shifts data in from MISO. Before a shift, the shift register is loaded with the contents of the SPI_TDBR register. At the end of the transfer, the contents of the shift register are loaded into SPI_RDBR.
6. With each new Transfer Initiate command, the SPI continues to send and receive words, according to the SPI Transfer Initiate mode.

If the transmit buffer remains empty or the receive buffer remains full, the device operates according to the states of the SZ and GM bits in SPI_CTL. If SZ = 1 and the transmit buffer is empty, the device repeatedly transmits 0s on the MOSI pin. One word is transmitted for each new Transfer Initiate command. If SZ = 0 and the transmit buffer is empty, the device

repeatedly transmits the last word it transmitted before the transmit buffer became empty. If GM = 1 and the receive buffer is full, the device continues to receive new data from the MISO pin, overwriting the older data in the SPI_RDBR buffer. If GM = 0 and the receive buffer is full, the incoming data is discarded, and SPI_RDBR is not updated.

Transfer Initiation From Master (Transfer Modes)

When a device is enabled as a master, the initiation of a transfer is defined by the two TIMOD bits of SPI_CTL. Based on those two bits and the status of the interface, a new transfer is started upon either a read of SPI_RDBR or a write to SPI_TDBR. This is summarized in Table 10-4.

Table 10-4. Transfer Initiation

TIMOD	Function	Transfer Initiated Upon	Action, Interrupt
00	Transmit and Receive	Initiate new single word transfer upon read of SPI_RDBR and previous transfer completed.	Interrupt active when receive buffer is full. Read of SPI_RDBR clears interrupt.
01	Transmit and Receive	Initiate new single word transfer upon write to SPI_TDBR and previous transfer completed.	Interrupt active when transmit buffer is empty. Writing to SPI_TDBR clears interrupt.
10	Receive with DMA	Initiate new multiword transfer upon enabling SPI for DMA mode. Individual word transfers begin with a DMA read of SPI_RDBR, and last transfer completed.	Interrupt active when DMA FIFO is not empty. Interrupt clears when DMA FIFO is empty.
11	Transmit with DMA	Initiate new multiword transfer upon enabling SPI for DMA mode. Individual word transfers begin with a DMA write to SPI_TDBR, and last transfer completed.	Interrupt active when DMA FIFO is not full. Interrupt clears when DMA FIFO is full.

Slave Mode Operation

When a device is enabled as a slave (and DMA mode is not selected), the start of a transfer is triggered by a transition of the $\overline{\text{SPISS}}$ select signal to the active state (low), or by the first active edge of the clock (SCK), depending on the state of CPHA.

These steps illustrate SPI operation in the Slave mode:

1. The core writes to SPI_CTL to define the mode of the serial link to be the same as the mode setup in the SPI master.
2. To prepare for the data transfer, the core writes data to be transmitted into SPI_TDBR.
3. Once the $\overline{\text{SPISS}}$ falling edge is detected, the slave starts shifting data out on MISO and in from MOSI on SCK edges, depending on the states of CPHA and CPOL.
4. Reception/transmission continues until $\overline{\text{SPISS}}$ is released or until the slave has received the proper number of clock cycles.
5. The slave device continues to receive/transmit with each new falling edge transition on $\overline{\text{SPISS}}$ and/or SCK clock edge.

If the transmit buffer remains empty or the receive buffer remains full, the device operates according to the states of the SZ and GM bits in SPI_CTL. If SZ = 1 and the transmit buffer is empty, the device repeatedly transmits 0s on the MISO pin. If SZ = 0 and the transmit buffer is empty, it repeatedly transmits the last word it transmitted before the transmit buffer became empty. If GM = 1 and the receive buffer is full, the device continues to receive new data from the MOSI pin, overwriting the older data in SPI_RDBR. If GM = 0 and the receive buffer is full, the incoming data is discarded, and SPI_RDBR is not updated.

Slave Ready for a Transfer

When a device is enabled as a slave, the actions shown in Table 10-5 are necessary to prepare the device for a new transfer.

Table 10-5. Transfer Preparation

TIMOD	Function	Action, Interrupt
00	Transmit and Receive	Interrupt active when receive buffer is full. Read of SPI_RDBR clears interrupt.
01	Transmit and Receive	Interrupt active when transmit buffer is empty. Writing to SPI_TDBR clears interrupt.
10	Receive with DMA	Interrupt active when DMA FIFO is not empty. Interrupt clears when DMA FIFO is empty.
11	Transmit with DMA	Interrupt active when DMA FIFO is not full. Interrupt clears when DMA FIFO is full.

Error Signals and Flags

The status of a device is indicated by the SPI_STAT register. See "SPI Status Register (SPI_STAT)" on page 10-15 for more information.

Mode Fault Error (MODF)

The MODF bit is set in SPI_STAT when the $\overline{\text{SPISS}}$ input pin of a device enabled as a master is driven low by some other device in the system. This occurs in multimaster systems when another device is also trying to be the master. To enable this feature, the PSSE bit in SPI_CTL must be set. This

contention between two drivers can potentially damage the driving pins. As soon as this error is detected, these actions occur:

- The MSTR control bit in SPI_CTL is cleared, configuring the SPI interface as a slave
- The SPE control bit in SPI_CTL is cleared, disabling the SPI system
- The MODF status bit in SPI_STAT is set
- An SPI Error interrupt is generated

These four conditions persist until the MODF bit is cleared by software. Until the MODF bit is cleared, the SPI cannot be re-enabled, even as a slave. Hardware prevents the user from setting either SPE or MSTR while MODF is set.

When MODF is cleared, the interrupt is deactivated. Before attempting to re-enable the SPI as a master, the state of the $\overline{\text{SPISS}}$ input pin should be checked to make sure the pin is high. Otherwise, once SPE and MSTR are set, another Mode Fault Error condition immediately occurs.

When SPE and MSTR are cleared, the SPI data and clock pin drivers (MOSI, MISO, and SCK) are disabled. However, the slave select output pins revert to being controlled by the programmable flag registers. This could lead to contention on the slave select lines if these lines are still driven by the processor. To ensure that the slave select output drivers are disabled once an MODF error occurs, the program must configure the programmable flag registers appropriately.

When enabling the MODF feature, the program must configure as inputs all of the PFx pins that will be used as slave selects. Programs can do this by configuring the direction of the PFx pins prior to configuring the SPI. This ensures that, once the MODF error occurs and the slave selects are automatically reconfigured as PFx pins, the slave select output drivers are disabled.

Transmission Error (TXE)

The TXE bit is set in SPI_STAT when all the conditions of transmission are met, and there is no new data in SPI_TDBR (SPI_TDBR is empty). In this case, the contents of the transmission depend on the state of the SZ bit in SPI_CTL. The TXE bit is sticky (W1C).

Reception Error (RBSY)

The RBSY flag is set in the SPI_STAT register when a new transfer is completed, but before the previous data can be read from SPI_RDBR. The state of the GM bit in the SPI_CTL register determines whether SPI_RDBR is updated with the newly received data. The RBSY bit is sticky (W1C).

Transmit Collision Error (TXCOL)

The TXCOL flag is set in SPI_STAT when a write to SPI_TDBR coincides with the load of the shift register. The write to SPI_TDBR can be via software or the DMA. The TXCOL bit indicates that corrupt data may have been loaded into the shift register and transmitted. In this case, the data in SPI_TDBR may not match what was transmitted. This error can easily be avoided by proper software control. The TXCOL bit is sticky (W1C).

Beginning and Ending an SPI Transfer

The start and finish of an SPI transfer depend on whether the device is configured as a master or a slave, whether the CPHA mode is selected, and whether the Transfer Initiation mode (TIMOD) is selected. For a master SPI with CPHA = 0, a transfer starts when either SPI_TDBR is written to or SPI_RDBR is read, depending on TIMOD. At the start of the transfer, the enabled slave select outputs are driven active (low). However, the SCK signal remains inactive for the first half of the first cycle of SCK. For a slave with CPHA = 0, the transfer starts as soon as the $\overline{\text{SPISS}}$ input goes low.

For CPHA = 1, a transfer starts with the first active edge of SCK for both slave and master devices. For a master device, a transfer is considered finished after it sends the last data and simultaneously receives the last data bit. A transfer for a slave device ends after the last sampling edge of SCK.

The RXS bit defines when the receive buffer can be read. The TXS bit defines when the transmit buffer can be filled. The end of a single word transfer occurs when the RXS bit is set, indicating that a new word has just been received and latched into the receive buffer, SPI_RDBR. For a master SPI, RXS is set shortly after the last sampling edge of SCK. For a slave SPI, RXS is set shortly after the last SCK edge, regardless of CPHA or CPOL. The latency is typically a few SCLK cycles and is independent of TIMOD and the baud rate. If configured to generate an interrupt when SPI_RDBR is full (TIMOD = 00), the interrupt goes active one SCLK cycle after RXS is set. When not relying on this interrupt, the end of a transfer can be detected by polling the RXS bit.

To maintain software compatibility with other SPI devices, the SPIF bit is also available for polling. This bit may have a slightly different behavior from that of other commercially available devices. For a slave device, SPIF is cleared shortly after the start of a transfer ($\overline{\text{SPISS}}$ going low for CPHA = 0, first active edge of SCK on CPHA = 1), and is set at the same time as RXS. For a master device, SPIF is cleared shortly after the start of a transfer (either by writing the SPI_TDBR or reading the SPI_RDBR, depending on TIMOD), and is set one-half SCK period after the last SCK edge, regardless of CPHA or CPOL.

The time at which SPIF is set depends on the baud rate. In general, SPIF is set after RXS, but at the lowest baud rate settings (SPI_BAUD < 4). The SPIF bit is set before RXS is set, and consequently before new data is latched into SPI_RDBR, because of the latency. Therefore, for SPI_BAUD = 2 or SPI_BAUD = 3, RXS must be set before SPIF to read SPI_RDBR. For larger SPI_BAUD settings, RXS is guaranteed to be set before SPIF is set.

If the SPI port is used to transmit and receive at the same time, or to switch between receive and transmit operation frequently, then the TIMOD = 00 mode may be the best operation option. In this mode, software performs a dummy read from the SPI_RDBR register to initiate the first transfer. If the first transfer is used for data transmission, software should write the value to be transmitted into the SPI_TDBR register before performing the dummy read. If the transmitted value is arbitrary, it is good practice to set the SZ bit to ensure zero data is transmitted rather than random values. When receiving the last word of an SPI stream, software should ensure that the read from the SPI_RDBR register does not initiate another transfer. It is recommended to disable the SPI port before the final SPI_RDBR read access. Reading the SPI_SHADOW register is not sufficient as it does not clear the interrupt request.

In master mode with the CPHA bit set, software should manually assert the required slave select signal before starting the transaction. After all data has been transferred, software typically releases the slave select again. If the SPI slave device requires the slave select line to be asserted for the complete transfer, this can be done in the SPI interrupt service routine only when operating in TIMOD = 00 or TIMOD = 10 mode. With TIMOD = 01 or TIMOD = 11, the interrupt is requested while the transfer is still in progress.

DMA

The SPI port also can use Direct Memory Access (DMA). For more information on DMA, see "Direct Memory Access" on page 9-1.

DMA Functionality

The SPI has a single DMA engine which can be configured to support either an SPI transmit channel or a receive channel, but not both simultaneously. Therefore, when configured as a transmit channel, the received data will essentially be ignored. When configured as a receive channel, what is transmitted is irrelevant. A 16-bit by four-word FIFO (without

burst capability) is included to improve throughput on the DMA Access Bus (DAB).

When using DMA for SPI transmit, the DMA_DONE interrupt signifies that the DMA FIFO is empty. However, at this point there may still be data in the SPI DMA FIFO waiting to be transmitted. Therefore, software needs to poll TXS in the SPI_STAT register until it goes low for 2 successive reads, at which point the SPI DMA FIFO will be empty. When the SPIF bit subsequently gets set, the last word has been transferred and the SPI can be disabled or enabled for another mode.

The four-word FIFO is cleared when the SPI port is disabled.

Master Mode DMA Operation

When enabled as a master with the DMA engine configured to transmit or receive data, the SPI interface operates as follows.

1. The processor core writes to the appropriate DMA registers to enable the SPI DMA Channel and to configure the necessary work units, access direction, word count, and so on. For more information, see "Direct Memory Access" on page 9-1.
2. The processor core writes to the SPI_FLG register, setting one or more of the SPI flag select bits (FLSx).
3. The processor core writes to the SPI_BAUD and SPI_CTL registers, enabling the device as a master and configuring the SPI system by specifying the appropriate word length, transfer format, baud rate, and so on. The TIMOD field should be configured to select either "Receive with DMA" (TIMOD = 10) or "Transmit with DMA" (TIMOD = 11) mode.

4. If configured for receive, a receive transfer is initiated upon enabling of the SPI. Subsequent transfers are initiated as the SPI reads data from the SPI_RDBR register and writes to the SPI DMA FIFO. The SPI then requests a DMA write to memory. Upon a DMA grant, the DMA engine reads a word from the SPI DMA FIFO and writes to memory.

   If configured for transmit, the SPI requests a DMA read from memory. Upon a DMA grant, the DMA engine reads a word from memory and writes to the SPI DMA FIFO. As the SPI writes data from the SPI DMA FIFO into the SPI_TDBR register, it initiates a transfer on the SPI link.

5. The SPI then generates the programmed clock pulses on SCK and simultaneously shifts data out of MOSI and shifts data in from MISO. For receive transfers, the value in the shift register is loaded into the SPI_RDBR register at the end of the transfer. For transmit transfers, the value in the SPI_TDBR register is loaded into the shift register at the start of the transfer.

6. In Receive mode, as long as there is data in the SPI DMA FIFO (the FIFO is not empty), the SPI continues to request a DMA write to memory. The DMA engine continues to read a word from the SPI DMA FIFO and writes to memory until the SPI DMA Word Count register transitions from 1 to 0. The SPI continues receiving words until SPI DMA mode is disabled.

   In Transmit mode, as long as there is room in the SPI DMA FIFO (the FIFO is not full), the SPI continues to request a DMA read from memory. The DMA engine continues to read a word from memory and write to the SPI DMA FIFO until the SPI DMA Word Count register transitions from 1 to 0. The SPI continues transmitting words until the SPI DMA FIFO is empty.

For receive DMA operations, if the DMA engine is unable to keep up with the receive datastream, the receive buffer operates according to the state of

the GM bit. If GM = 1 and the DMA FIFO is full, the device continues to receive new data from the MISO pin, overwriting the older data in the SPI_RDBR register. If GM = 0, and the DMA FIFO is full, the incoming data is discarded, and the SPI_RDBR register is not updated. While performing receive DMA, the transmit buffer is assumed to be empty (and TXE is set). If SZ = 1, the device repeatedly transmits 0s on the MOSI pin. If SZ = 0, it repeatedly transmits the contents of the SPI_TDBR register. The TXE underrun condition cannot generate an error interrupt in this mode.

For transmit DMA operations, the master SPI initiates a word transfer only when there is data in the DMA FIFO. If the DMA FIFO is empty, the SPI waits for the DMA engine to write to the DMA FIFO before starting the transfer. All aspects of SPI receive operation should be ignored when configured in Transmit DMA mode, including the data in the SPI_RDBR register, and the status of the RXS and RBSY bits. The RBSY overrun conditions cannot generate an error interrupt in this mode. The TXE underrun condition cannot happen in this mode (master DMA TX mode), because the master SPI will not initiate a transfer if there is no data in the DMA FIFO.

Writes to the SPI_TDBR register during an active SPI transmit DMA operation should not occur because the DMA data will be overwritten. Writes to the SPI_TDBR register during an active SPI receive DMA operation are allowed. Reads from the SPI_RDBR register are allowed at any time.

Data interrupts (DMA requests) are generated when the DMA FIFO is not empty (when TIMOD = 10), or when the DMA FIFO is not full (when TIMOD = 11).

Error interrupts are generated when there is an RBSY overflow error condition (when TIMOD = 10).

A master SPI DMA sequence may involve back-to-back transmission and/or reception of multiple DMA work units. The SPI controller supports such a sequence with minimal core interaction.

Slave Mode DMA Operation

When enabled as a slave with the DMA engine configured to transmit or receive data, the start of a transfer is triggered by a transition of the $\overline{\text{SPISS}}$ signal to the active-low state or by the first active edge of SCK, depending on the state of CPHA.

The following steps illustrate the SPI receive or transmit DMA sequence in an SPI slave (in response to a master command).

1. The processor core writes to the appropriate DMA registers to enable the SPI DMA Channel and configure the necessary work units, access direction, word count, and so on. For more information, see "Direct Memory Access" on page 9-1.

2. The processor core writes to the SPI_CTL register to define the mode of the serial link to be the same as the mode setup in the SPI master. The TIMOD field will be configured to select either "Receive with DMA" (TIMOD = 10) or "Transmit with DMA" (TIMOD = 11) mode.

3. If configured for receive, once the slave select input is active, the slave starts receiving and transmitting data on SCK edges. The value in the shift register is loaded into the SPI_RDBR register at the end of the transfer. As the SPI reads data from the SPI_RDBR register and writes to the SPI DMA FIFO, it requests a DMA write to memory. Upon a DMA grant, the DMA engine reads a word from the SPI DMA FIFO and writes to memory.

   If configured for transmit, the SPI requests a DMA read from memory. Upon a DMA grant, the DMA engine reads a word from memory and writes to the SPI DMA FIFO. The SPI then reads data from the SPI DMA FIFO and writes to the SPI_TDBR register, awaiting the start of the next transfer. Once the slave select input is

active, the slave starts receiving and transmitting data on SCK edges. The value in the SPI_TDBR register is loaded into the shift register at the start of the transfer.

4. In Receive mode, as long as there is data in the SPI DMA FIFO (FIFO not empty), the SPI slave continues to request a DMA write to memory. The DMA engine continues to read a word from the SPI DMA FIFO and writes to memory until the SPI DMA Word Count register transitions from 1 to 0. The SPI slave continues receiving words on SCK edges as long as the slave select input is active.

   In Transmit mode, as long as there is room in the SPI DMA FIFO (FIFO not full), the SPI slave continues to request a DMA read from memory. The DMA engine continues to read a word from memory and write to the SPI DMA FIFO until the SPI DMA Word Count register transitions from 1 to 0. The SPI slave continues transmitting words on SCK edges as long as the slave select input is active.

For receive DMA operations, if the DMA engine is unable to keep up with the receive datastream, the receive buffer operates according to the state of the GM bit. If GM = 1 and the DMA FIFO is full, the device continues to receive new data from the MOSI pin, overwriting the older data in the SPI_RDBR register. If GM = 0 and the DMA FIFO is full, the incoming data is discarded, and the SPI_RDBR register is not updated. While performing receive DMA, the transmit buffer is assumed to be empty and TXE is set. If SZ = 1, the device repeatedly transmits 0s on the MISO pin. If SZ = 0, it repeatedly transmits the contents of the SPI_TDBR register. The TXE underrun condition cannot generate an error interrupt in this mode.

For transmit DMA operations, if the DMA engine is unable to keep up with the transmit stream, the transmit port operates according to the state of the SZ bit. If SZ = 1 and the DMA FIFO is empty, the device repeatedly transmits 0s on the MISO pin. If SZ = 0 and the DMA FIFO is empty, it repeatedly transmits the last word it transmitted before the

DMA buffer became empty. All aspects of SPI receive operation should be ignored when configured in Transmit DMA mode, including the data in the SPI_RDBR register, and the status of the RXS and RBSY bits. The RBSY overrun conditions cannot generate an error interrupt in this mode.

Writes to the SPI_TDBR register during an active SPI transmit DMA operation should not occur because the DMA data will be overwritten. Writes to the SPI_TDBR register during an active SPI receive DMA operation are allowed. Reads from the SPI_RDBR register are allowed at any time.

Data interrupts (DMA requests) are generated when the DMA FIFO is not empty (when TIMOD = 10), or when the DMA FIFO is not full (when TIMOD = 11).

Error interrupts are generated when there is an RBSY overflow error condition (when TIMOD = 10), or when there is a TXE underflow error condition (when TIMOD = 11).

Timing

The enable lead time (T1), the enable lag time (T2), and the sequential transfer delay time (T3) each must always be greater than or equal to one-half the SCK period. See Figure 10-14. The minimum time between successive word transfers (T4) is two SCK periods. This is measured from the last active edge of SCK of one word to the first active edge of SCK of the next word. This is independent of the configuration of the SPI (CPHA, MSTR, and so on).

For a master device with CPHA = 0, the slave select output will be inactive (high) for at least one-half the SCK period. In this case, T1 and T2 will each always be equal to one-half the SCK period.

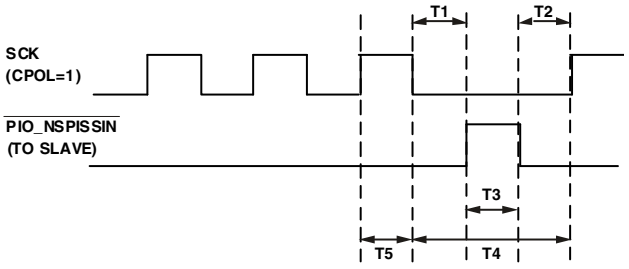

Figure 10-14. SPI Timing

11 PARALLEL PERIPHERAL INTERFACE

The ADSP-BF561 provides two 16-bit Parallel Peripheral Interfaces (PPI0 and PPI1) that can connect directly to parallel A/D and D/A converters, video encoders and decoders and other general purpose peripherals.

The Parallel Peripheral Interface (PPI) is a half-duplex, bidirectional port accommodating up to 16 bits of data. It has a dedicated clock pin, three frame sync pins, and eight dedicated data pins.

The default behavior is for frame syncs and data to be sampled on opposite PPI clock edges. To change this behavior such that frame sync and data are sampled on the same PPI clock edge, set the ALT_TIMING bit in the PLL_CTL register. From more information about PPI controls, see Figure 11-1.

Up to eight additional data pins are available by reconfiguring the PF pins. The highest system throughput is achieved with 8-bit data, since two 8-bit data samples can be packed as a single 16-bit word. In such a case, the earlier sample is placed in the 8 least significant bits (LSBs).

The PPIx_CLK pin can accept an external clock input. It cannot source a clock internally. Table 11-1 shows the pin interface for the PPI.

If a programmable flag pin is configured for PPI use, its bit position in programmable flag MMRs will read back as 0.

Table 11-1. PPIx Pins (x=0, 1, Represents PPI0 and PPI1)

Signal Name	Function	Direction	Alternate Function
PPIx_15	Data	Bidirectional	Programmable Flag Pin
PPIx_14	Data	Bidirectional	Programmable Flag Pin
PPIx_13	Data	Bidirectional	Programmable Flag Pin
PPIx_12	Data	Bidirectional	Programmable Flag Pin
PPIx_11	Data	Bidirectional	Programmable Flag Pin
PPIx_10	Data	Bidirectional	Programmable Flag Pin
PPIx_9	Data	Bidirectional	Programmable Flag Pin
PPIx_8	Data	Bidirectional	Programmable Flag Pin
PPIx_7	Data	Bidirectional	N/A
PPIx_6	Data	Bidirectional	N/A
PPIx_5	Data	Bidirectional	N/A
PPIx_4	Data	Bidirectional	N/A
PPIx_3	Data	Bidirectional	N/A
PPIx_2	Data	Bidirectional	N/A
PPIx_1	Data	Bidirectional	N/A
PPIx_0	Data	Bidirectional	N/A
PPIx_FS3	Frame Sync3/Field	Bidirectional	N/A
PPIx_FS2	Frame Sync2/VSYNC	Bidirectional	Timer
PPIx_FS1	Frame Sync1/HSYNC	Bidirectional	Timer
PPIx_CLK		Input Clock	N/A

PPI Memory-Mapped Registers

The PPI has five memory-mapped registers (MMRs) that regulate its operation. These registers are the PPI Control register (PPIx_CONTROL), the PPI Status register (PPIx_STATUS), the Delay Count register (PPIx_DELAY), the Transfer Count register (PPIx_COUNT), and the Lines Per Frame register (PPIx_FRAME).

Note upon disabling the PPI, each of these MMRs will be reset to its default value. To insure intended operation, the appropriate MMRs may need to be written before enabling the PPI again.

Descriptions and bit diagrams for each of these MMRs are provided in the following sections.

PPI Control Register (PPIx_CONTROL)

This register configures the PPI for operating mode, control signal polarities, and data width of the port. See Figure 11-1 for a bit diagram of this MMR.

The POLC and POLS bits allow for selective signal inversion of the PPIx_CLK and PPIx_FS1/PPIx_FS2 signals, respectively. This provides a mechanism to connect to data sources and receivers with a wide array of control signal polarities. Often, the remote data source/receiver also offers configurable signal polarities, so the POLC and POLS bits simply add increased flexibility.

The DLEN[2:0] field is programmed to specify the width of the PPI port in any mode. Note any width from 8 to 16 bits is supported, with the exception of a 9-bit port width. Any PF pins that are unused by the PPI as a result of the DLEN setting are free to be used in their normal PF capacity.

In ITU-R 656 modes, the DLEN field should not be configured for anything greater than a 10-bit port width. If it is, the PPI will reserve extra pins, making them unusable by other peripherals.

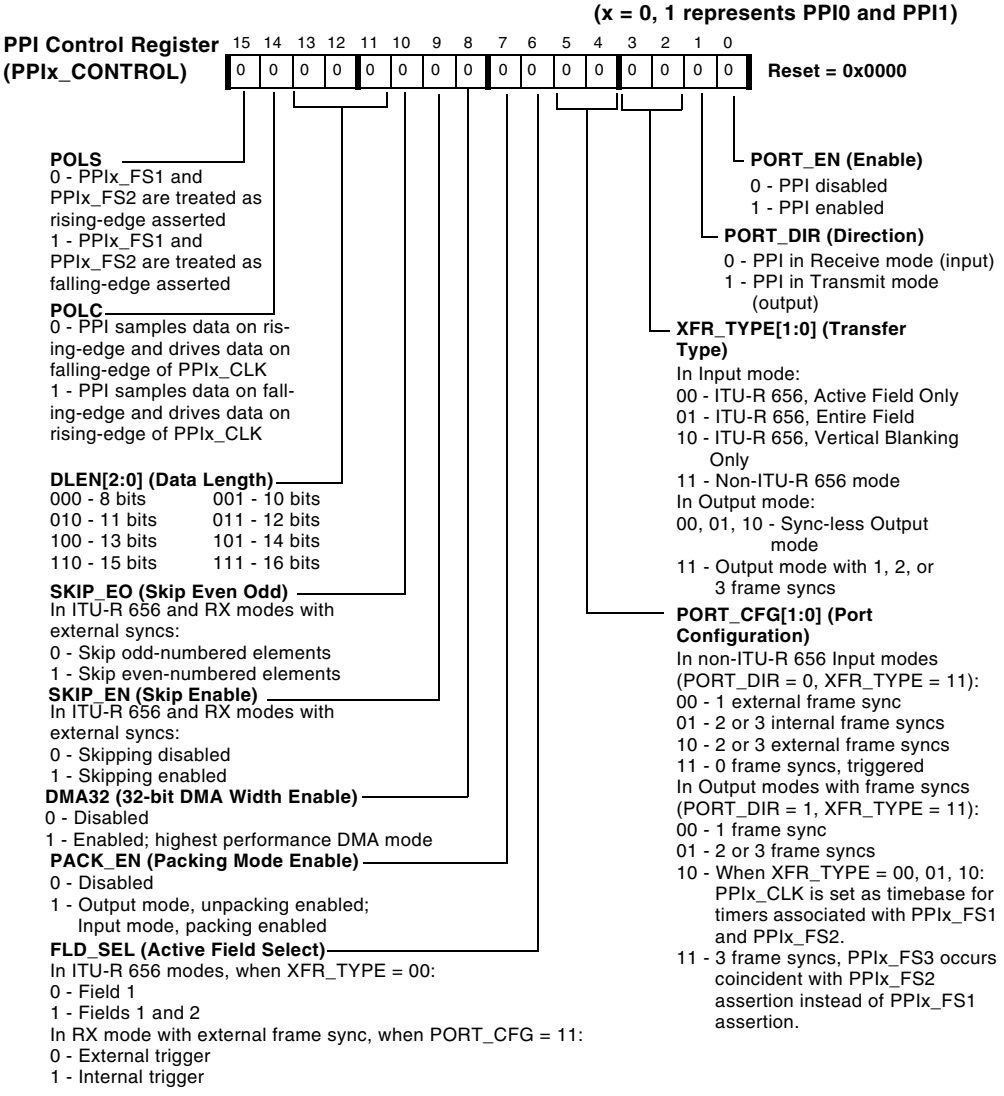

Figure 11-1. PPI Control Register

The SKIP_EN bit, when set, enables the selective skipping of data elements being read in through the PPI. By ignoring data elements, the PPI is able to conserve DMA bandwidth.

When the SKIP_EN bit is set, the SKIP_EO bit allows the PPI to ignore either the odd or the even elements in an input datastream. This is useful, for instance, when reading in a color video signal in YCbCr format (Cb, Y, Cr, Y, Cb, Y, Cr, Y...). Skipping every other element allows the PPI to only read in the luma (Y) or chroma (Cr or Cb) values. This could also be useful when synchronizing two PPI's on the same or different processors to the same incoming video stream. One PPI could handle luma processing and the other (whose SKIP_EO bit is set differently from the first) could handle chroma processing. This skipping feature is valid in ITU-R 656 modes and RX modes with external frame syncs.

The PACK_EN bit allows data packing/unpacking to/from 32-bit words. Every PPIx_CLK-initiated event on the DMA bus (that is, an input or output operation) handles 32-bit entities. In other words, an input port width of 10 bits still results in a 32-bit input word for every PPIx_CLK; the upper 22 bits are 0s. Likewise, a port width of 8 bits also results in a 32-bit input word, with the upper 24 bits all 0s. In the case of 8-bit data, it is usually more efficient to pack this information so that there are four bytes of data for every 32-bit word. This is the function of the PACK_EN bit. When set, it enables packing for all RX modes. For TX modes, setting PACK_EN enables unpacking of bytes.

The DMA32 bit allows for different packing options when used in conjunction with the PACK_EN. Refer to Table 11-2 and Table 11-3.

Table 11-2. Packing Options (Data Sent)

Data sent from memory over DMA bus to PPI: 0x4433 2211, 0x8877 6655, 0xCCBB AA99, 0x11FF EEDD			
DLEN	**PACK_EN**	**DMA32**	**Output from PPI**
>8 bits	0	0	0x2211, 0x6655, 0xAA99, 0xEEDD
8 bits	0	0	0x0011, 0x0055, 0x0099, 0x00DD
>8 bits	0	1	0x2211, 0x4433, 0x6655, 0x8877, 0xAA99, 0xCCBB, 0xEEDD, 0x11FF
8 bits	0	1	0x0011, 0x0033, 0x0055, 0x0077, 0x0099, 0x00BB, 0x00DD, 0x00FF
>8 bits	1	0	0x0011, 0x0022, 0x0055, 0x0066, 0x0099, 0x00AA, 0x00DD, 0x00EE
8 bits	1	0	0x0011, 0x0022, 0x0055, 0x0066, 0x0099, 0x00AA, 0x00DD, 0x00EE
>8 bits	1	1	0x0011, 0x0022, 0x0033, 0x0044 ,0x0055, 0x0066, 0x0077, 0x0088, 0x0099, 0x00AA, 0x00BB, 0x00CC, 0x00DD, 0x00EE, 0x00FF. 0x0011
8 bits	1	1	0x0011, 0x0022, 0x0033, 0x0044 ,0x0055, 0x0066, 0x0077, 0x0088, 0x0099, 0x00AA, 0x00BB, 0x00CC, 0x00DD, 0x00EE, 0x00FF. 0x0011

Table 11-3. Packing Options (Data Received)

Data received into PPI: 0x2211, 0x4433, 0x6655, 0x8877			
DLEN	**PACK_EN**	**DMA32**	**Input to PPI**
>8 bits	0	0	0x2211, 0x4433, 0x6655, 0x8877
8 bits	0	0	0x0011, 0x0033, 0x0055, 0x0077
>8 bits	0	1	0x4433 2211, 0x8877 6655
8 bits	0	1	0x0033 0011, 0x0055 0077

Table 11-3. Packing Options (Data Received)

Data received into PPI: 0x2211, 0x4433, 0x6655, 0x8877			
DLEN	**PACK_EN**	**DMA32**	**Input to PPI**
>8 bits	1	0	0x2211, 0x4433, 0x6655, 0x8877
8 bits	1	0	0x3311, 0x7755
>8 bits	1	1	0x4433 2211, 0x8877 6655
8 bits	1	1	0x7755 3311

The FLD_SEL bit is used primarily in the Active Field Only ITU-R 656 mode. The FLD_SEL bit determines whether to transfer in only Field 1 of each video frame, or both Fields 1 and 2. Thus, it allows a savings in DMA bandwidth by transferring only every other field of active video.

The PORT_CFG[1:0] field is used to configure the operating mode of the PPI. It operates in conjunction with the PORT_DIR bit, which sets the direction of data transfer for the port. The XFR_TYPE[1:0] field is also used to configure operating mode and is discussed below. See Table 11-4 for the possible operating modes for the PPI.

PORT_CFG = 10 in TX mode, 0 frame syncs sets PPIx_CLK as the timebase for the timers associated with PPIx_FS1 and PPIx_FS2.

Table 11-4. PPI Possible Operating Modes

PPI Mode	Number of Syncs	PORT_DIR	PORT_CFG	XFR_TYPE	POLC	POLS	FLD_ SEL
RX mode, 0 frame syncs, external trigger	0	0	11	11	0 or 1	0 or 1	0
RX mode, 0 frame syncs, internal trigger	0	0	11	11	0 or 1	0 or 1	1
RX mode, 1 external frame sync	1	0	00	11	0 or 1	0 or 1	X
RX mode, 2 or 3 external frame syncs	3	0	10	11	0 or 1	0 or 1	X
RX mode, 2 or 3 internal frame syncs	3	0	01	11	0 or 1	0 or 1	X
RX mode, ITU-R 656, Active Field Only	embedded	0	XX	00	0 or 1	0	0 or 1
RX mode, ITU-R 656, Vertical Blanking Only	embedded	0	XX	10	0 or 1	0	X
RX mode, ITU-R 656, Entire Field	embedded	0	XX	01	0 or 1	0	X
TX mode, 0 frame syncs	0	1	XX	00, 01, 10	0 or 1	0 or 1	X
TX mode, 1 internal or external frame sync	1	1	00	11	0 or 1	0 or 1	X
TX mode, 2 external frame syncs	2	1	01	11	0 or 1	0 or 1	X

Table 11-4. PPI Possible Operating Modes (Cont'd)

PPI Mode	Number of Syncs	PORT_DIR	PORT_CFG	XFR_TYPE	POLC	POLS	FLD_ SEL
TX mode, 2 or 3 internal frame syncs, FS3 sync'd to FS1 assertion	3	1	01	11	0 or 1	0 or 1	X
TX mode, 2 or 3 internal frame syncs, FS3 sync'd to FS2 assertion	3	1	11	11	0 or 1	0 or 1	X

The XFR_TYPE[1:0] field configures the PPI for various modes of operation. Refer to Table 11-4 to see how XFR_TYPE[1:0] interacts with other bits in PPIx_CONTROL to determine the PPI operating mode.

The PORT_EN bit, when set, enables the PPI for operation.

When configured as an input port, the PPI does not start data transfer after being enabled until the appropriate synchronization signals are received. If configured as an output port, transfer (including the appropriate synchronization signals) begins as soon as the frame syncs (Timer units) are enabled, so all frame syncs must be configured before this happens. Refer to the section "Frame Synchronization in GP Modes" on page 11-29 for more information.

PPI Status Register (PPIx_STATUS)

The PPI Status register (Figure 11-2) contains bits that provide information about the current operating state of the PPI.

The entire register is cleared when read, so the status word must be parsed to evaluate which bits have been set.

The ERR_DET bit is a sticky bit that denotes whether or not an error was detected in the ITU-R 656 control word preamble. The bit is valid only in ITU-R 656 modes. If ERR_DET = 1, an error was detected in the preamble. If ERR_DET = 0, no error was detected in the preamble.

The ERR_NCOR bit is sticky and is relevant only in ITU-R 656 modes. If ERR_NCOR = 0, an error in the preamble was detected and corrected. If ERR_NCOR = 1, an error in the preamble was detected but not corrected. This situation generates a PPI Error interrupt, unless this condition is masked off in the SIC_IMASK0 register.

The FT_ERR bit is sticky and indicates, when set, that a Frame Track Error has occurred. It is valid for RX modes only. In this condition, the programmed number of lines per frame in PPIx_FRAME does not match up with the "frame start detect" condition (see the information note on page 11-13). A Frame Track Error generates a PPI Error interrupt, unless this condition is masked off in the SIC_IMASK0 register.

The FLD bit is set or cleared at the same time as the change in state of F (in ITU-R 656 modes) or PPIx_FS3 (in other RX modes). It is valid for Input modes only. The state of FLD reflects the current state of the F or PPIx_FS3 signals. In other words, the FLD bit always reflects the current video field being processed by the PPI.

The OVR bit is sticky and indicates, when set, that the PPI FIFO has overflowed and can accept no more data. A FIFO Overflow Error generates a PPI Error interrupt, unless this condition is masked off in the SIC_IMASK0 register.

The PPI FIFO is 32 bits wide and has 16 entries.

The UNDR bit is sticky and indicates, when set, that the PPI FIFO has underrun and is data-starved. A FIFO Underrun Error generates a PPI Error interrupt, unless this condition is masked off in the SIC_IMASK0 register.

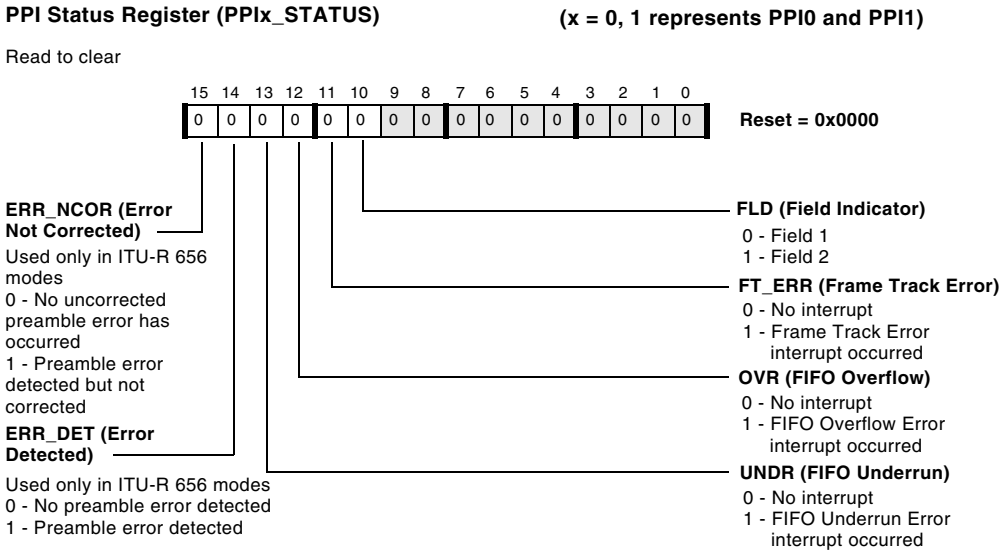

Figure 11-2. PPI Status Register

Delay Count Register (PPIx_DELAY)

The PPIx_DELAY register (Figure 11-3) can be used in all configurations except ITU-R 656 modes and GP modes with 0 frame syncs. It contains a count of how many PPIx_CLK cycles to delay after assertion of PPIx_FS1 before starting to read in or write out data.

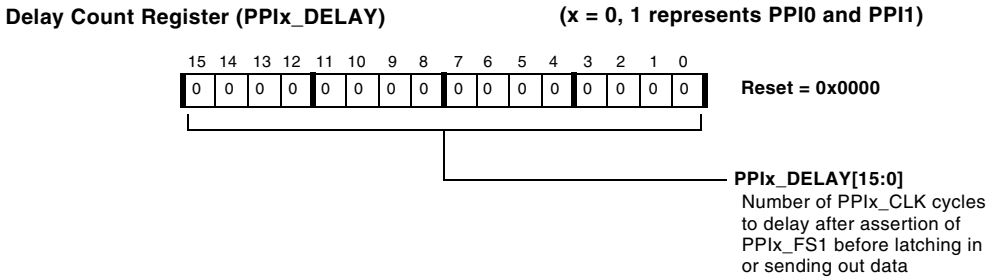

Figure 11-3. Delay Count Register

Note in TX modes using at least one frame sync, there is a one-cycle delay beyond what is specified in the PPIx_DELAY register.

Transfer Count Register (PPIx_COUNT)

The PPIx_COUNT register (Figure 11-4) is used only in cases where recurring hardware frame syncs (either externally or internally generated) are involved. It is not needed in ITU-R 656 modes or modes with zero frame syncs. For RX modes, this register holds the number of samples to read into the PPI per line, minus one. For TX modes, it holds the number of samples to write out through the PPI per line, minus one. The register itself does not actually decrement with each transfer. Thus, at the beginning of a new line of data, there is no need to rewrite the value of this register. For example, to receive or transmit 100 samples through the PPI, set PPIx_COUNT to 99.

Take care to ensure that the number of samples programmed into PPIx_COUNT is in keeping with the number of samples expected during the "horizontal" interval specified by PPIx_FS1.

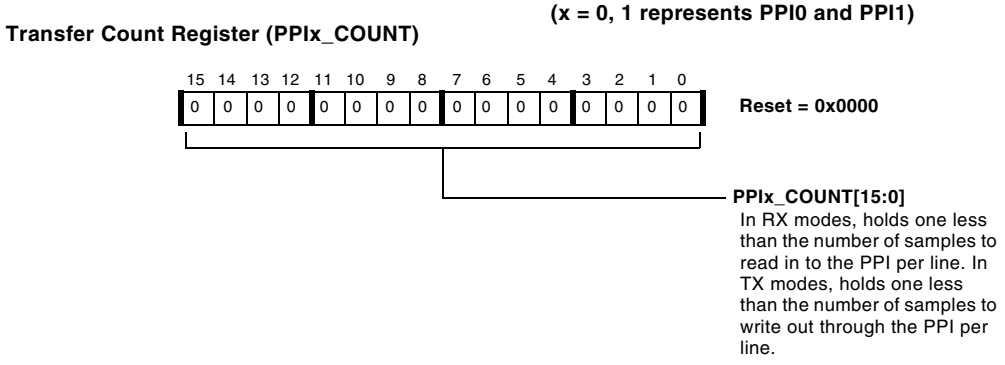

Figure 11-4. Transfer Count Register

Lines Per Frame Register (PPIx_FRAME)

The Lines Per Frame (PPIx_FRAME) register (Figure 11-5) is used in TX mode with 2 or 3 external frame syncs, as well as in all RX modes with 2 or 3 frame syncs. For ITU-R 656 modes, this register holds the number of lines expected per frame of data, where a frame is defined as Field 1 and Field 2 combined, designated by the “F” indicator in the ITU-R stream. Here, a line is defined as a complete ITU-R 656 SAV-EAV cycle.

For non-ITU-R 656 modes with external frame syncs, a frame is defined as the data bounded between PPIx_FS2 assertions, regardless of the state of PPIx_FS3. A line is defined as a complete PPIx_FS1 cycle. In these modes, PPIx_FS3 is used only to determine the original "frame start" each time PPIx is enabled. It is ignored on every subsequent field and frame, and its state (high or low) is not important except during the original frame start.

If the start of a new frame (or field, for ITU-R 656 mode) is detected before the number of lines specified by PPIx_FRAME have been transferred, a Frame Track Error results, and the FT_ERR bit in PPIx_STATUS is set. However, PPIx still automatically reinitializes to count to the value programmed in PPIx_FRAME, and data transfer continues.

In ITU-R 656 modes, a frame start detect happens on the falling edge of F, the Field indicator. This occurs at the start of Field 1.

In RX mode with 3 external frame syncs, a frame start detect refers to a condition where a PPIx_FS2 assertion is followed by an assertion of PPIx_FS1 while PPIx_FS3 is low. This occurs at the start of Field 1. Note that PPIx_FS3 only needs to be low when PPIx_FS1 is asserted, not when PPIx_FS2 asserts. Also, PPIx_FS3 is only used to synchronize to the start of the very first frame after PPIx is enabled. It is subsequently ignored.

When using RX mode with 3 external frame syncs, and only 2 syncs are needed, configure the PPI for three-frame-sync operation and provide an external pull-down to GND for the PPIx_FS3 pin.

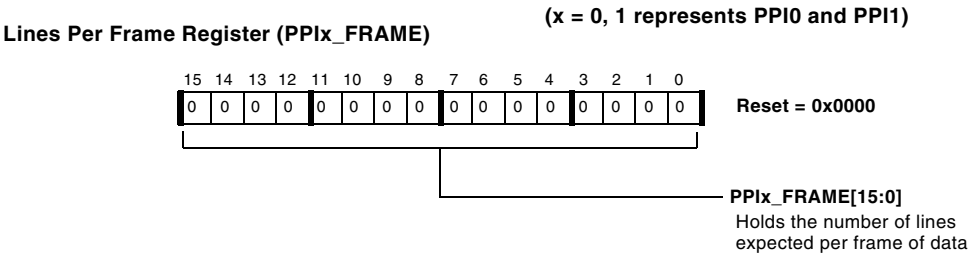

Figure 11-5. Lines Per Frame Register

ITU-R 656 Modes

The PPI supports three input modes for ITU-R 656-framed data. These modes are described in this section. Although the PPI does not explicitly support an ITU-R 656 output mode, recommendations for using the PPI for this situation are provided as well.

ITU-R 656 Background

According to the ITU-R 656 recommendation (formerly known as CCIR-656), a digital video stream has the characteristics shown in Figure 11-6 and Figure 11-7 for 525/60 (NTSC) and 625/50 (PAL) systems. The processor supports only the Bit-parallel mode of ITU-R 656. Both 8- and 10-bit video element widths are supported.

In this mode, the Horizontal (H), Vertical (V), and Field (F) signals are sent as an embedded part of the video datastream in a series of bytes that form a control word. The Start of Active Video (SAV) and End of Active Video (EAV) signals indicate the beginning and end of data elements to read in on each line. SAV occurs on a 1-to-0 transition of H, and EAV begins on a 0-to-1 transition of H. An entire field of video is comprised of Active Video + Horizontal Blanking (the space between an EAV and SAV code) and Vertical Blanking (the space where V = 1). A field of video

commences on a transition of the F bit. The “odd field” is denoted by a value of F = 0, whereas F = 1 denotes an even field. Progressive video makes no distinction between Field 1 and Field 2, whereas interlaced video requires each field to be handled uniquely, because alternate rows of each field combine to create the actual video image.

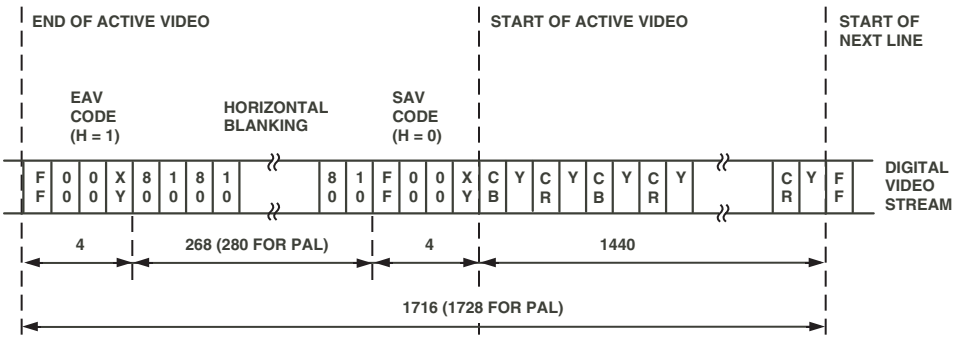

Figure 11-6. ITU-R 656 8-Bit Parallel Data Stream for NTSC (PAL) Systems

The SAV and EAV codes are shown in more detail in Table 11-5. Note there is a defined preamble of three bytes (0xFF, 0x00, 0x00), followed by the XY Status word, which, aside from the F (Field), V (Vertical Blanking) and H (Horizontal Blanking) bits, contains four protection bits for single-bit error detection and correction. Note F and V are only allowed to change as part of EAV sequences (that is, transition from H = 0 to H = 1). The bit definitions are as follows:

- F = 0 for Field 1
- F = 1 for Field 2
- V = 1 during Vertical Blanking
- V = 0 when not in Vertical Blanking

- H = 0 at SAV
- H = 1 at EAV
- P3 = V XOR H
- P2 = F XOR H
- P1 = F XOR V
- P0 = F XOR V XOR H

In many applications, video streams other than the standard NTSC/PAL formats (for example, CIF, QCIF) can be employed. Because of this, the processor interface is flexible enough to accommodate different row and field lengths. In general, as long as the incoming video has the proper EAV/SAV codes, the PPI can read it in. In other words, a CIF image could be formatted to be "656-compliant," where EAV and SAV values define the range of the image for each line, and the V and F codes can be used to delimit fields and frames.

Table 11-5. Control Byte Sequences for 8- and 10-Bit ITU-R 656 Video

	8-bit Data								**10-bit Data**	
	D9 (MSB)	D8	D7	D6	D5	D4	D3	D2	D1	D0
Preamble	1	1	1	1	1	1	1	1	1	1
	0	0	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	0	0	0
Control Byte	1	F	V	H	P3	P2	P1	P0	0	0

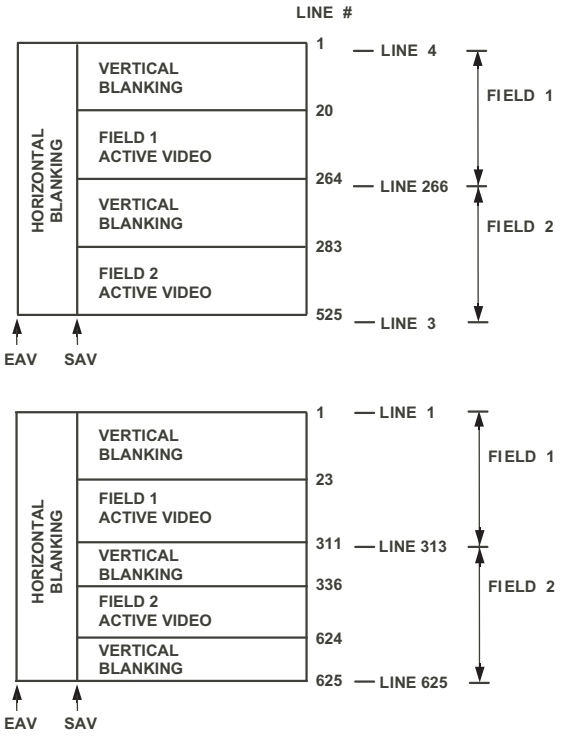

LINE NUMBER	F	V	H (EAV)	H (SAV)
1-3, 266-282	1	1	1	0
4-19, 264-265	0	1	1	0
20-263	0	0	1	0
283-525	1	0	1	0

LINE NUMBER	F	V	H (EAV)	H (SAV)
1-22, 311-312	0	1	1	0
23-310	0	0	1	0
313-335, 624-625	1	1	1	0
336-623	1	0	1	0

Figure 11-7. Typical Video Frame Partitioning for NTSC/PAL Systems for ITU-R BT.656-4

ITU-R 656 Input Modes

Figure 11-8 shows a general illustration of data movement in the ITU-R 656 input modes. In the figure, the clock `CLK` is either provided by the video source or supplied externally by the system.

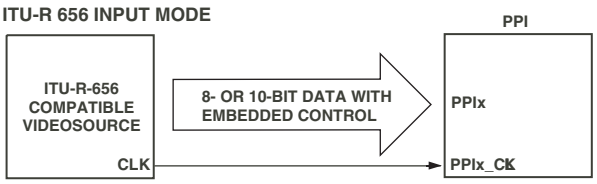

Figure 11-8. ITU-R 656 Input Modes

There are three submodes supported for ITU-R 656 inputs: Entire Field, Active Video Only, and Vertical Blanking Interval Only. Figure 11-9 shows these three submodes.

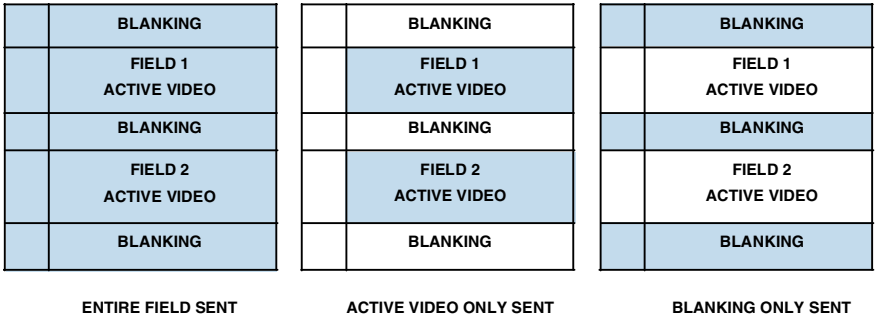

Figure 11-9. ITU-R 656 Input Submodes

Entire Field

In this mode, the entire incoming bitstream is read in through the PPI. This includes Active Video as well as control byte sequences and ancillary data that may be embedded in Horizontal and Vertical Blanking Intervals. Data transfer starts immediately after synchronization to Field 1 occurs, but does not include the first EAV code that contains the F = 0 assignment.

Note the first line transferred in after enabling the PPI will be missing its first 4-byte preamble. However, subsequent lines and frames should have all control codes intact.

One side benefit of this mode is that it enables a “loopback” feature through which a frame or two of data can be read in through the PPI and subsequently output to a compatible video display device. Of course, this requires multiplexing on the PPI pins, but it enables a convenient way to verify that ITU-R 656 data can be read into and written out from the PPI.

Active Video Only

This mode is used when only the active video portion of a field is of interest, and not any of the blanking intervals. The PPI ignores (does not read in) all data between EAV and SAV, as well as all data present when V = 1. In this mode, the control byte sequences are not stored to memory; they are filtered out by the PPI. After synchronizing to the start of Field 1, the PPI ignores incoming samples until it sees an SAV.

In this mode, the user specifies the number of total (active plus vertical blanking) lines per frame in the `PPIx_FRAME` MMR.

Vertical Blanking Interval (VBI) Only

In this mode, data transfer is only active while V = 1 is in the control byte sequence. This indicates that the video source is in the midst of the Vertical Blanking Interval (VBI), which is sometimes used for ancillary data

transmission. The ITU-R 656 recommendation specifies the format for these ancillary data packets, but the PPI is not equipped to decode the packets themselves. This task must be handled in software. Horizontal blanking data is logged where it coincides with the rows of the VBI. Control byte sequence information is always logged. The user specifies the number of total lines (Active plus Vertical Blanking) per frame in the PPIx_FRAME MMR.

Note the VBI is split into two regions within each field. From the PPI's standpoint, it considers these two separate regions as one contiguous space. However, keep in mind that frame synchronization begins at the start of Field 1, which doesn't necessarily correspond to the start of Vertical Blanking. For instance, in 525/60 systems, the start of Field 1 (F = 0) corresponds to Line 4 of the VBI.

ITU-R 656 Output Mode

The PPI does not explicitly provide functionality for framing an ITU-R 656 output stream with proper preambles and blanking intervals. However, with the TX mode with 0 frame syncs, this process can be supported manually. Essentially, this mode provides a streaming operation from memory out through the PPI. Data and control codes can be set up in memory prior to sending out the video stream. With the 2D DMA engine, this could be performed in a number of ways. For instance, one line of blanking (H + V) could be stored in a buffer and sent out N times by the DMA controller when appropriate, before proceeding to DMA active video. Alternatively, one entire field (with control codes and blanking) can be set up statically in a buffer while the DMA engine transfers only the active video region into the buffer, on a frame-by-frame basis.

Frame Synchronization in ITU-R 656 Modes

Synchronization in ITU-R 656 modes always occurs at the falling edge of F, the field indicator. This corresponds to the start of Field 1.

Consequently, up to two fields might be ignored (for example, if Field 1 just started before the PPI-to-camera channel was established) before data is received into the PPI.

Because all H and V signalling is embedded in the datastream in ITU-R 656 modes, the PPIx_COUNT register is not necessary. However, the PPIx_FRAME register is used in order to check for synchronization errors. The user programs this MMR for the number of lines expected in each frame of video, and the PPI keeps track of the number of EAV-to-SAV transitions that occur from the start of a frame until it decodes the end-of-frame condition (transition from F = 1 to F = 0). At this time, the actual number of lines processed is compared against the value in PPIx_FRAME. If there is a mismatch, the FT_ERR bit in the PPIx_STATUS register is asserted. For instance, if an SAV transition is missed, the current field will only have NUM_ROWS - 1 rows, but resynchronization will reoccur at the start of the next frame.

Upon completing reception of an entire field, the Field Status bit is toggled in the PPIx_STATUS register. This way, an interrupt service routine (ISR) can discern which field was just read in.

General-Purpose PPI Modes

The General-Purpose (GP) PPI modes are intended to suit a wide variety of data capture and transmission applications. Table 11-6 summarizes these modes. If a particular mode shows a given PPI_FSx frame sync not being used, this implies that the pin is available for its alternate, multiplexed processor function (that is, as a timer or flag pin). The exception to this is that when a PPI is configured for a 2-frame-sync mode, PPI_FS3 cannot be used as a general purpose flag, even though it is not used by the PPI.

Table 11-6. General-Purpose PPI Modes

GP PPI Mode	PPIx_FS1 Direction	PPIx_FS2 Direction	PPIx_FS3 Direction	Data Direction
RX mode, 0 frame syncs, external trigger	Input	Not used	Not used	Input
RX mode, 0 frame syncs, internal trigger	Not used	Not used	Not used	Input
RX mode, 1 external frame sync	Input	Not used	Not used	Input
RX mode, 2 or 3 external frame syncs	Input	Input	Input	Input
RX mode, 2 or 3 internal frame syncs	Output	Output	Output	Input
TX mode, 0 frame syncs	Not used	Not used	Not used	Output
TX mode, 1 external frame sync	Input	Not used	Not used	Output
TX mode, 2 external frame syncs	Input	Input	Output	Output
TX mode, 1 internal frame sync	Output	Not used	Output	Output
TX mode, 2 or 3 internal frame syncs	Output	Output	Output	Output

Figure 11-10 illustrates the general flow of the GP modes. The top of the diagram shows an example of RX mode with 1 external frame sync. After the PPI receives the hardware frame sync pulse (PPIx_FS1), it delays for the duration of the PPIx_CLK cycles programmed into PPIx_DELAY. The DMA controller then transfers in the number of samples specified by PPIx_COUNT. Every sample that arrives after this, but before the next PPIx_FS1 frame sync arrives, is ignored and not transferred onto the DMA bus.

If the next PPIx_FS1 frame sync arrives before the specified PPIx_COUNT samples have been read in, the sample counter reinitializes to 0 and starts to count up to PPIx_COUNT again. This situation can cause the DMA channel configuration to lose synchronization with the PPI transfer process.

The bottom of Figure 11-10 shows an example of TX mode, 1 internal frame sync. After PPIx_FS1 is asserted, there is a latency of 1 PPIx_CLK cycle, and then there is a delay for the number of PPIx_CLK cycles

programmed into PPIx_DELAY. Next, the DMA controller transfers out the number of samples specified by PPIx_COUNT. No further DMA takes place until the next PPIx_FS1 sync and programmed delay occur.

If the next PPIx_FS1 frame sync arrives before the specified PPIx_COUNT samples have been transferred out, the sync has priority and starts a new line transfer sequence. This situation can cause the DMA channel configuration to lose synchronization with the PPI transfer process.

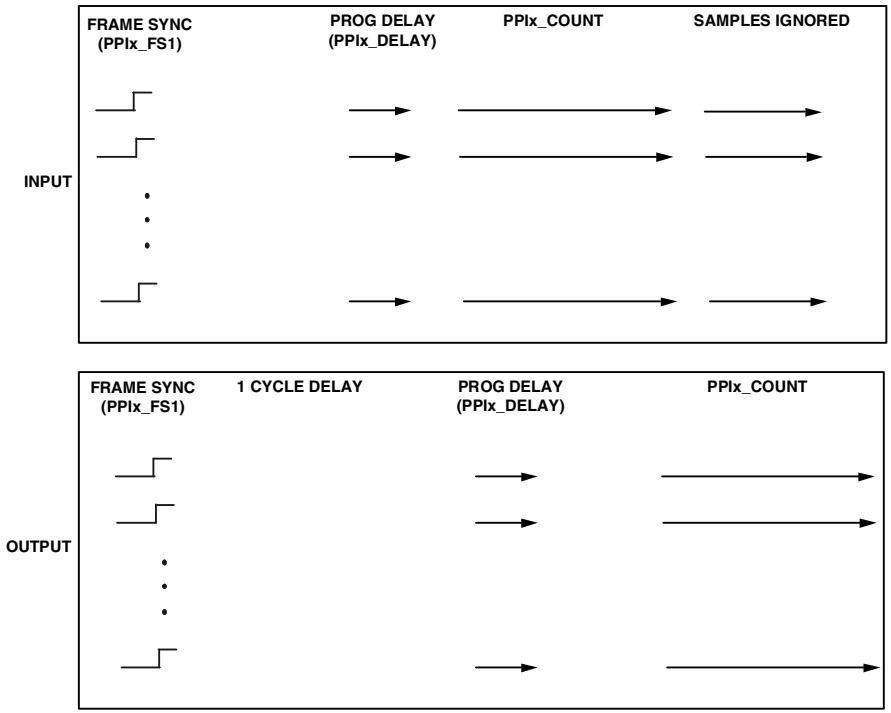

Figure 11-10. General Flow for GP Modes (Assumes Positive Assertion of PPIx_FS1)

Data Input (RX) Modes

The PPI supports several modes for data input. These modes differ chiefly by the way the data is framed. Refer to Table 11-4 on page 11-8 for information on how to configure the PPI for each mode.

No Frame Syncs

These modes cover the set of applications where periodic frame syncs are not generated to frame the incoming data. There are two options for starting the data transfer, both configured by the PPIx_CONTROL register.

- External trigger: An external source sends a single frame sync (tied to PPIx_FS1) at the start of the transaction, when FLD_SEL = 0 and PORT_CFG = b#11.
- Internal trigger: Software initiates the process by setting PORT_EN = 1 with FLD_SEL = 1 and PORT_CFG = b#11.

All subsequent data manipulation is handled via DMA. For example, an arrangement could be set up between alternating 1K memory buffers. When one fills up, DMA continues with the second buffer, at the same time that another DMA operation is clearing the first memory buffer for reuse.

Due to clock domain synchronization in RX modes with no frame syncs, there may be a delay of at least two PPI_CLK cycles between when the mode is enabled and when valid data is received. Therefore, detection of the start of valid data should be managed by software.

1, 2, or 3 External Frame Syncs

The 1-sync mode is intended for analog-to-digital converter (ADC) applications. The top part of Figure 11-11 shows a typical illustration of the system setup for this mode.

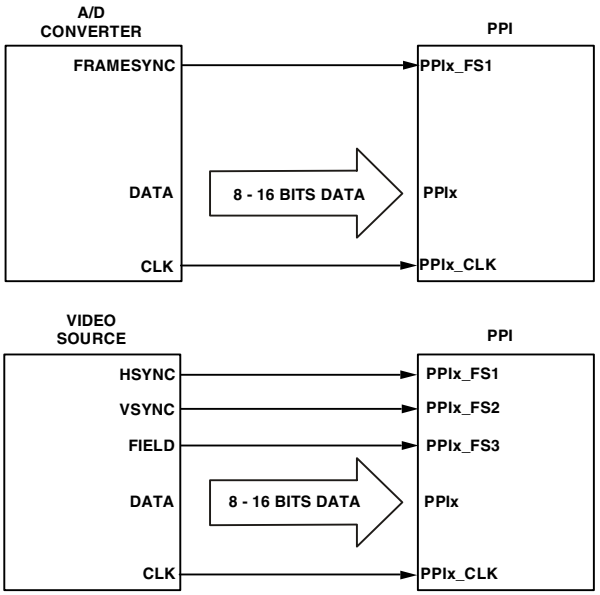

Figure 11-11. RX Mode, External Frame Syncs

The 3-sync mode shown at the bottom of Figure 11-11 supports video applications that use hardware signalling (HSYNC, VSYNC, FIELD) in accordance with the ITU-R 601 recommendation. The mapping for the frame syncs in this mode is PPIx_FS1 = HSYNC, PPIx_FS2 = VSYNC, PPIx_FS3 = FIELD. Please refer to "Frame Synchronization in GP Modes" on page 11-29 for more information about frame syncs in this mode.

A 2-sync mode is implicitly supported by pulling PPIx_FS3 to GND by an external resistor when configured in 3-sync mode.

2 or 3 Internal Frame Syncs

This mode can be useful for interfacing to video sources that can be slaved to a master processor. In other words, the processor controls when to read from the video source by asserting PPIx_FS1 and PPIx_FS2, and then reading data into the PPI. The PPIx_FS3 frame sync provides an indication of which field is currently being transferred, but since it is an output, it can simply be left floating if not used. Figure 11-12 shows a sample application for this mode.

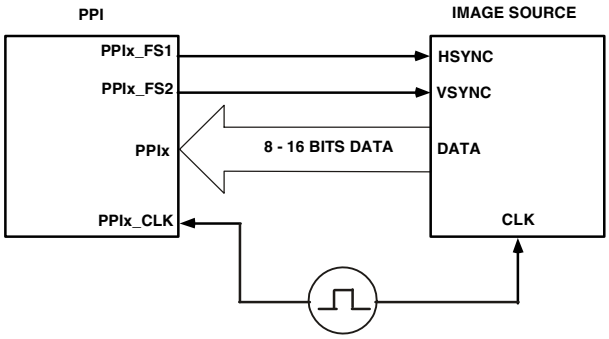

Figure 11-12. RX Mode, Internal Frame Syncs

Data Output (TX) Modes

The PPI supports several modes for data output. These modes differ chiefly by the way the data is framed. Refer to Table 11-4 on page 11-8 for information on how to configure the PPI for each mode.

No Frame Syncs

In this mode, data blocks specified by the DMA controller are sent out through the PPI with no framing. That is, once the DMA channel is configured and enabled, and the PPI is configured and enabled, data transfers

will take place immediately, synchronized to PPIx_CLK. See Figure 11-13 for an illustration of this mode.

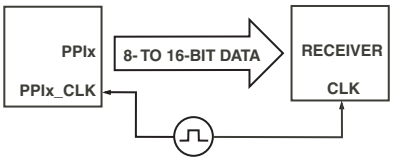

Figure 11-13. TX Mode, 0 Frame Syncs

In this mode, there is a one-time, minimum delay of 16 SCLK cycles (for > 8-bit data) or 32 SCLK cycles (for 8-bit data) between enabling the PPI and transmission of valid data. Furthermore, DMA must be configured to transmit at least 16 samples (for > 8-bit data) or 32 samples (for 8-bit data).

1 or 2 External Frame Syncs

In these modes, an external receiver can frame data sent from the PPI. Both 1-sync and 2-sync modes are supported. The top diagram in Figure 11-14 shows the 1-sync case, while the bottom diagram illustrates the 2-sync mode.

There is a mandatory delay of 1.5 PPIx_CLK cycles, plus the value programmed in PPIx_DELAY, between assertion of the external frame sync(s) and the transfer of valid data out through the PPI.

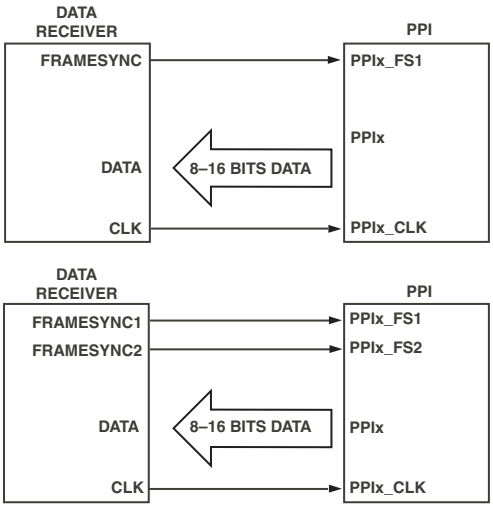

Figure 11-14. TX Mode, 1 or 2 External Frame Syncs

1, 2, or 3 Internal Frame Syncs

The 1-sync mode is intended for interfacing to digital-to-analog converters (DACs) with a single frame sync. The top part of Figure 11-15 shows an example of this type of connection.

The 3-sync mode is useful for connecting to video and graphics displays, as shown in the bottom part of Figure 11-15. A 2-sync mode is implicitly supported by leaving `PPIx_FS3` unconnected in this case.

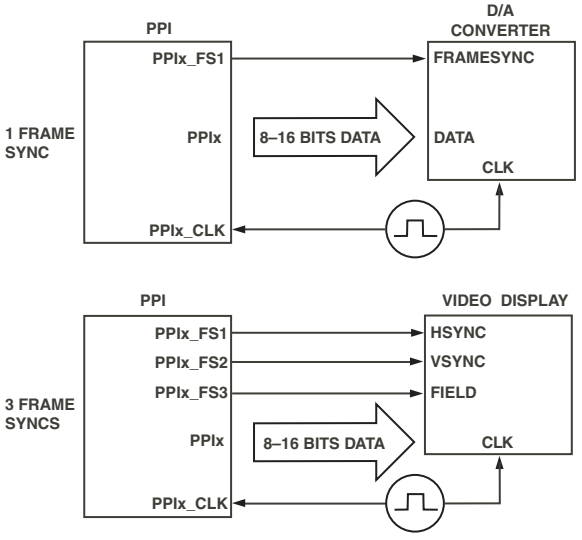

Figure 11-15. TX Mode, 1–3 Internal Frame Syncs

Frame Synchronization in GP Modes

Frame synchronization in GP modes operates differently in modes with internal frame syncs than in modes with external frame syncs.

Modes With Internal Frame Syncs

In modes with internal frame syncs, PPIx_FS1 and PPIx_FS2 link directly to the Pulsewidth Modulation (PWM) circuits of Timers 8 and 9 (for PPI0) or Timers 10 and 11 (for PPI1). This allows for arbitrary pulse widths and periods to be programmed for these signals using the existing TIMERx registers. This capability accommodates a wide range of timing needs. Note these PWM circuits are clocked by PPIx_CLK. If PPIx_FS2 is not used in the configured PPI mode, the corresponding timer operates as it normally would, unrestricted in functionality. The state of PPIx_FS3

depends completely on the state of PPIx_FS1 and/or PPIx_FS2, so PPIx_FS3 has no inherent programmability.

To program PPIx_FS1 and/or PPIx_FS2 for operation in an internal frame sync mode:

1. Configure and enable DMA for the PPI. See "DMA Operation" on page 11-31.
2. Configure the width and period for PPIx_FS1 and PPIx_FS2 (if used) via TIMERx_WIDTH and TIMERx_PERIOD.
3. Set up TIMERx_CONFIG for PWM_OUT mode for PPIx_FS1. If used, configure TIMERx_CONFIG for PWM_OUT mode for PPIx_FS2. This includes setting CLK_SEL = 1 and TIN_SEL = 1 for each timer.
4. Write to PPIx_CONTROL to configure and enable the PPI.
5. Write to TMRS4_ENABLE to enable the appropriate timers.

It is important to guarantee proper frame sync polarity between the PPI and Timer peripherals. To do this, make sure that if PPIx_CONTROL[15:14] = b#10 or b#11, the PULSE_HI bit is cleared in the TIMERx_CONFIG registers corresponding to PPIx_FS1 and PPIx_FS2. Likewise, if PPIx_CONTROL[15:14] = b#00 or b#01, the PULSE_HI bit should be set in TIMER1_CONFIG and TIMER2_CONFIG.

To switch to another PPI mode not involving internal frame syncs:

1. Disable the PPI (using PPIx_CONTROL).
2. Disable the timers (using TMRS4_DISABLE).

Modes With External Frame Syncs

In RX modes with external frame syncs, the PPIx_FS1 and PPIx_FS2 pins become edge-sensitive inputs. In such a mode, the corresponding Timers can be used for a purpose not involving their respective TMRx pins.

However, timer access to a TMRx pin is disabled when the PPI is using that pin for a frame sync input function. For modes that do not require PPIx_FS2, the corresponding Timer is not restricted in functionality and can be operated as if the PPI were not being used (that is, that TMRx pin becomes available for timer use as well). For more information on configuring and using the timers, please refer to Chapter 15, "Timers".

In RX Mode with 3 external frame syncs, the start of frame detection occurs where a PPIx_FS2 assertion is followed by an assertion of PPIx_FS1 while PPIx_FS3 is low. This happens at the start of Field 1. Note that PPIx_FS3 only needs to be low when PPIx_FS1 is asserted, not when PPIx_FS2 asserts. Also, PPIx_FS3 is only used to synchronize to the start of the very first frame after PPIx is enabled. It is subsequently ignored.

In TX modes with external frame syncs, the PPIx_FS1 and PPIx_FS2 pins are also treated as edge-sensitive inputs. However, it is necessary to explicitly configure the timer(s) associated with the frame sync(s) as inputs, and to enable them. Therefore, write 0x0003 to the TIMERx_CONFIG registers associated with PPIx_FS1 and, if used, PPIx_FS2. This should be done before enabling the PPI. After the PPI has been enabled, write to TMRS4_ENABLE to enable the appropriate timers.

DMA Operation

The PPI must be used with the processor's DMA engine. This section discusses how the two interact. For additional information about the DMA engine, including explanations of DMA registers and DMA operations, please refer to Chapter 9, "Direct Memory Access".

Each PPI DMA channel can be configured for either transmit or receive operation, and it has a maximum throughput of (PPIx_CLK) × (16 bits/transfer). The highest throughput is achieved with 8-bit data, PACK_EN = 1 (packing mode enabled), and DMA32 = 1. Note for 32-bit

packing mode, the number of data elements to transfer must be a multiple of four.

Configuring the PPI's DMA channel is a necessary step toward using the PPI interface. It is the DMA engine that generates interrupts upon completion of a row, frame, or partial-frame transfer. It is also the DMA engine that coordinates the origination or destination point for the data that is transferred through the PPI.

The processor's 2D DMA capability allows the processor to be interrupted at the end of a line or after a frame of video has been transferred, as well as if a DMA Error occurs. In fact, the specification of the `DMAx_y_XCOUNT` and `DMAx_y_YCOUNT` MMRs allows for flexible data interrupt points. For example, assume the DMA registers `XMODIFY = YMODIFY = 1`. Then, if a data frame contains 320 × 240 bytes (240 rows of 320 bytes each), these conditions hold:

- Setting `XCOUNT = 320`, `YCOUNT = 240`, and `DI_SEL = 1` (the `DI_SEL` bit is located in `DMAx_y_CONFIG`) will interrupt on every row transferred, for the entire frame.
- Setting `XCOUNT = 320`, `YCOUNT = 240`, and `DI_SEL = 0` will interrupt only on the completion of the frame (when 240 rows of 320 bytes have been transferred).
- Setting `XCOUNT = 38,400` (320 × 120), `YCOUNT = 2`, and `DI_SEL = 1` will cause an interrupt when half of the frame has been transferred, and again when the whole frame has been transferred.

Following is the general procedure for setting up DMA operation with the PPI. Please refer to "Direct Memory Access" on page 9-1 for details regarding configuration of DMA.

1. Configure DMA registers as appropriate for desired DMA operating mode.
2. Enable the DMA channel for operation.

3. Configure appropriate PPI registers.
4. Enable the PPI by writing a `1` to bit 0 in `PPIx_CONTROL`.

Data Transfer Scenarios

Figure 11-16 shows two possible ways to use the PPI to transfer in video. These diagrams are very generalized, and bandwidth calculations must be made only after factoring in the exact PPI mode and settings (for example, transfer Field 1 only, transfer odd and even elements).

The top part of the diagram shows a situation appropriate for, as an example, JPEG compression. The first N rows of video are DMA'd into L1 memory via the PPI. Once in L1, the compression algorithm operates on the data and sends the compressed result out from the processor via the SPORT. Note that no SDRAM access was necessary in this approach.

The bottom part of the diagram takes into account a more formidable compression algorithm, such as MPEG-2 or MPEG-4. Here, the raw video is transferred directly into SDRAM. Independently, a Memory DMA channel transfers data blocks between SDRAM and L1 memory for intermediate processing stages. Finally, the compressed video exits the processor via the SPORT.

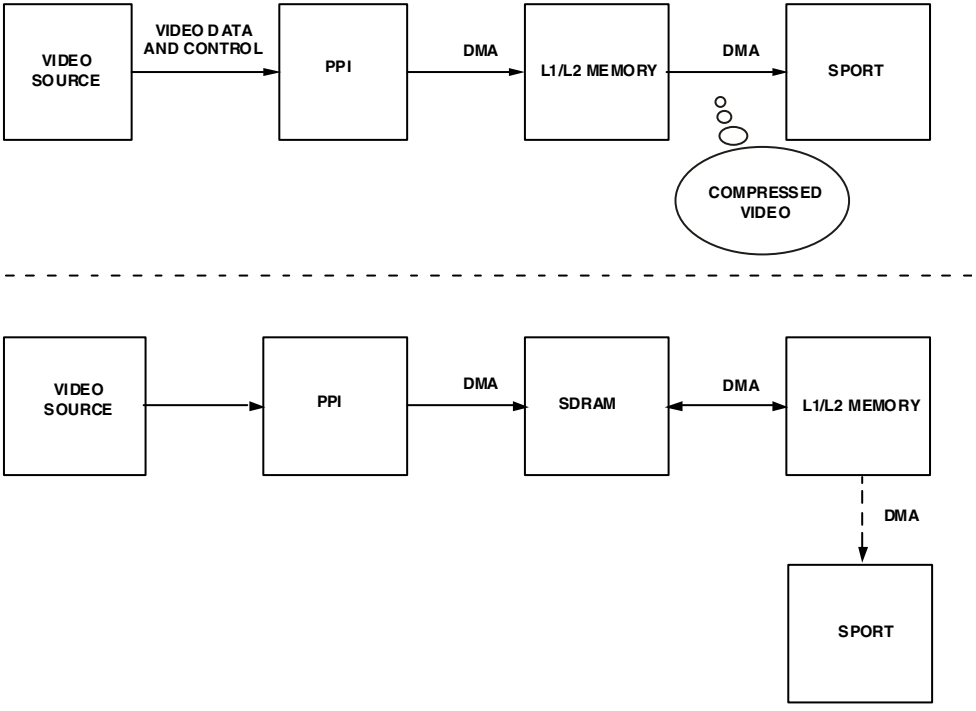

Figure 11-16. PPI Possible Data Transfer Scenarios

12 SERIAL PORT CONTROLLERS

The processor has two identical synchronous serial ports, or SPORTs. These support a variety of serial data communications protocols and can provide a direct interconnection between processors in a multiprocessor system. See Figure 12-1 on page 12-6.

The serial ports (SPORT0 and SPORT1) provide an I/O interface to a wide variety of peripheral serial devices. SPORTs provide synchronous serial data transfer only; the processor provides asynchronous RS-232 data transfer via the UART. Each SPORT has one group of pins (primary data, secondary data, clock, and frame sync) for transmit and a second set of pins for receive. The receive and transmit functions are programmed separately. Each SPORT is a full duplex device, capable of simultaneous data transfer in both directions. The SPORTs can be programmed for bit rate, frame sync, and number of bits per word by writing to memory-mapped registers.

In this text, the naming conventions for registers and pins use a lower case `x` to represent a digit. In this chapter, for example, the name `RFSx` pins indicates `RFS0` and `RFS1` (corresponding to SPORT0 and SPORT1, respectively). In this chapter, LSB refers to least significant bit, and MSB refers to most significant bit.

Both SPORTs have the same capabilities and are programmed in the same way. Each SPORT has its own set of control registers and data buffers.

The SPORTs use frame sync pulses to indicate the beginning of each word or packet, and the bit clock marks the beginning of each data bit. External bit clock and frame sync are available for the TX and RX buffers.

With a range of clock and frame synchronization options, the SPORTs allow a variety of serial communication protocols, including H.100, and provide a glueless hardware interface to many industry-standard data converters and codecs.

The SPORTs can operate at up to an SCLK/2 clock rate with an externally generated clock, or 1/2 the system clock rate for an internally generated serial port clock. The SPORT external clock must always be less than the SCLK frequency. Independent transmit and receive clocks provide greater flexibility for serial communications.

SPORT clocks and frame syncs can be internally generated by the system or received from an external source. The SPORTs can operate with a transmission format of LSB first or MSB first, with word lengths selectable from 3 to 32 bits. They offer selectable transmit modes and optional μ-law or A-law companding in hardware. SPORT data can be automatically transferred between on-chip and off-chip memories using DMA block transfers. Additionally, each of the SPORTs offers a TDM (Time-Division-Multiplexed) Multichannel mode.

Each of the SPORTs offers these features and capabilities:

- Provides independent transmit and receive functions.
- Transfers serial data words from 3 to 32 bits in length, either MSB first or LSB first.
- Provides alternate framing and control for interfacing to I^2S serial devices, as well as other audio formats (for example, left-justified stereo serial data).
- Has FIFO plus double buffered data (both receive and transmit functions have a data buffer register and a Shift register), providing additional time to service the SPORT.

- Provides two synchronous transmit and two synchronous receive data pins and buffers in each SPORT to double the total supported datastreams.
- Performs A-law and μ-law hardware companding on transmitted and received words. (See "Companding" on page 12-33 for more information.)
- Internally generates serial clock and frame sync signals in a wide range of frequencies or accepts clock and frame sync input from an external source.
- Operates with or without frame synchronization signals for each data word, with internally generated or externally generated frame signals, with active high or active low frame signals, and with either of two configurable pulse widths and frame signal timing.
- Performs interrupt-driven, single word transfers to and from on-chip memory under processor control.
- Provides Direct Memory Access transfer to and from memory under DMA Master control. DMA can be autobuffer-based (a repeated, identical range of transfers) or descriptor-based (individual or repeated ranges of transfers with differing DMA parameters).
- Executes DMA transfers to and from on-chip memory. Each SPORT can automatically receive and transmit an entire block of data.
- Permits chaining of DMA operations for multiple data blocks.
- Has a multichannel mode for TDM interfaces. Each SPORT can receive and transmit data selectively from a Time-Division-Multiplexed serial bitstream on 128 contiguous channels from a stream of up to 1024 total channels. This mode can be useful as a network communication scheme for multiple processors. The 128 channels available to the processor can be selected to start at any channel

location from 0 to 895 = (1023 – 128). Note the Multichannel Select registers and the WSIZE register control which subset of the 128 channels within the active region can be accessed.

Table 12-1 shows the pins for each SPORT.

Table 12-1. Serial Port (SPORT) Pins

Pin[1]	Description
DTxPRI	Transmit Data Primary
DTxSEC	Transmit Data Secondary
TSCLKx	Transmit Clock
TFSx	Transmit Frame Sync
DRxPRI	Receive Data Primary
DRxSEC	Receive Data Secondary
RSCLKx	Receive Clock
RFSx	Receive Frame Sync

1 A lowercase x within a pin name represents a possible value of 0 or 1 (corresponding to SPORT0 or SPORT1).

A SPORT receives serial data on its DRxPRI and DRxSEC inputs and transmits serial data on its DTxPRI and DTxSEC outputs. It can receive and transmit simultaneously for full-duplex operation. For transmit, the data bits (DTxPRI and DTxSEC) are synchronous to the transmit clock (TSCLKx). For receive, the data bits (DRxPRI and DRxSEC) are synchronous to the receive clock (RSCLKx). The serial clock is an output if the processor generates it, or an input if the clock is externally generated. Frame synchronization signals RFSx and TFSx are used to indicate the start of a serial data word or stream of serial words.

The primary and secondary data pins provide a method to increase the data throughput of the serial port. They do not behave as totally separate SPORTs; rather, they operate in a synchronous manner (sharing clock and

frame sync) but on separate data. The data received on the primary and secondary pins is interleaved in main memory and can be retrieved by setting a stride in the Data Address Generators (DAG) unit. For more information about DAGs, see Chapter 5, "Data Address Generators". Similarly, for TX, data should be written to the TX register in an alternating manner—first primary, then secondary, then primary, then secondary, and so on. This is easily accomplished with the processor's powerful DAGs.

In addition to the serial clock signal, data must be signalled by a frame synchronization signal. The framing signal can occur either at the beginning of an individual word or at the beginning of a block of words.

The following figure shows a simplified block diagram of a single SPORT. Data to be transmitted is written from an internal processor register to the SPORT's SPORTx_TX register via the peripheral bus. This data is optionally compressed by the hardware and automatically transferred to the TX Shift register. The bits in the Shift register are shifted out on the SPORT's DTxPRI/DTxSEC pin, MSB first or LSB first, synchronous to the serial clock on the TSCLKx pin. The receive portion of the SPORT accepts data from the DRxPRI/DRxSEC pin synchronous to the serial clock on the RSCLKx pin. When an entire word is received, the data is optionally expanded, then automatically transferred to the SPORT's SPORTx_RX register, and then into the RX FIFO where it is available to the processor.

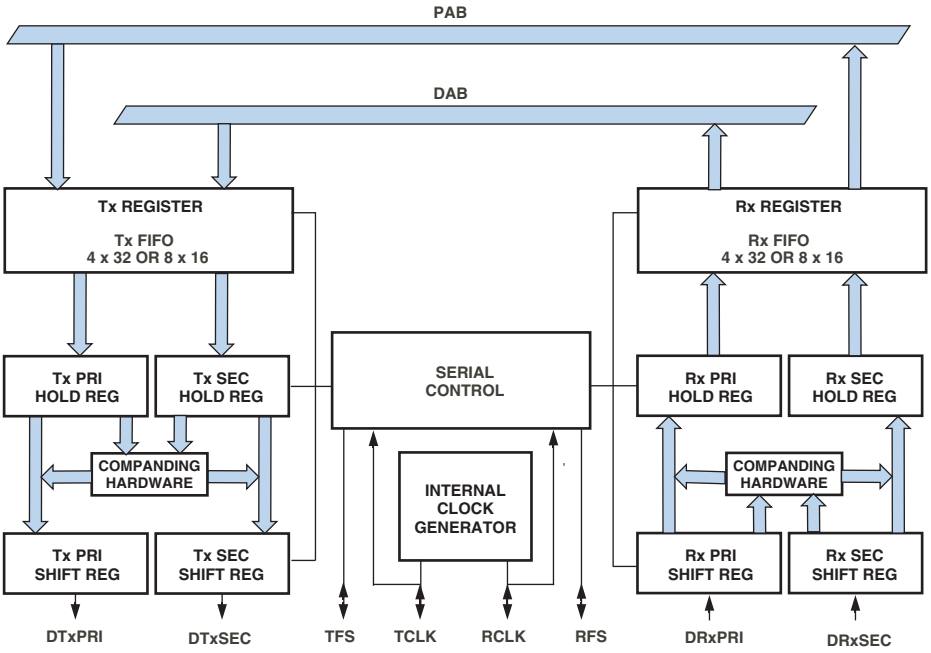

NOTE 1: ALL WIDE ARROW DATA PATHS ARE 16 OR 32 BITS WIDE, DEPENDING ON SLEN. FOR SLEN = 2 TO 15, A 16-BIT DATA PATH WITH 8-DEEP FIFO IS USED. FOR SLEN = 16 TO 31, A 32-BIT DATA PATH WITH 4-DEEP FIFO IS USED.
NOTE 2: Tx REGISTER IS THE BOTTOM OF THE Tx FIFO, Rx REGISTER IS THE TOP OF THE Rx FIFO.

Figure 12-1. SPORT Block Diagram

Figure 12-2 shows a possible port connection for the SPORTs. Note serial devices A and B must be synchronous, as they share common frame syncs and clocks. The same is true for serial devices C and D.

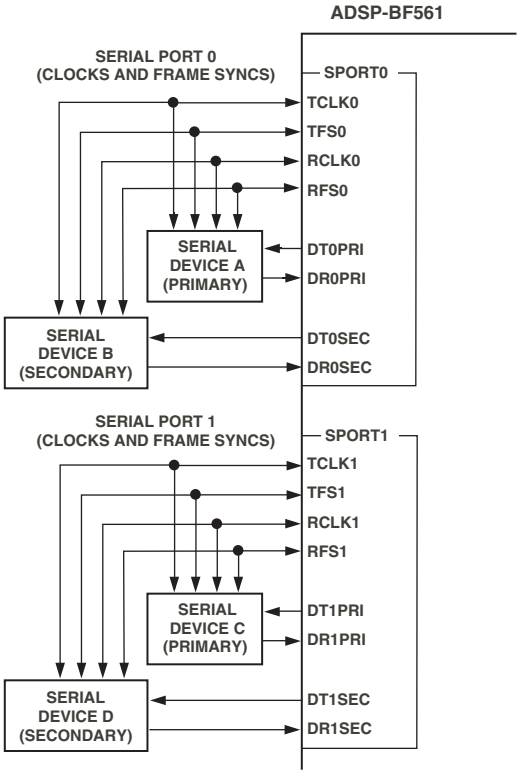

Figure 12-2. SPORT Connections

Figure 12-3 shows an example of a stereo serial device with three transmit and two receive channels connected to the processor.

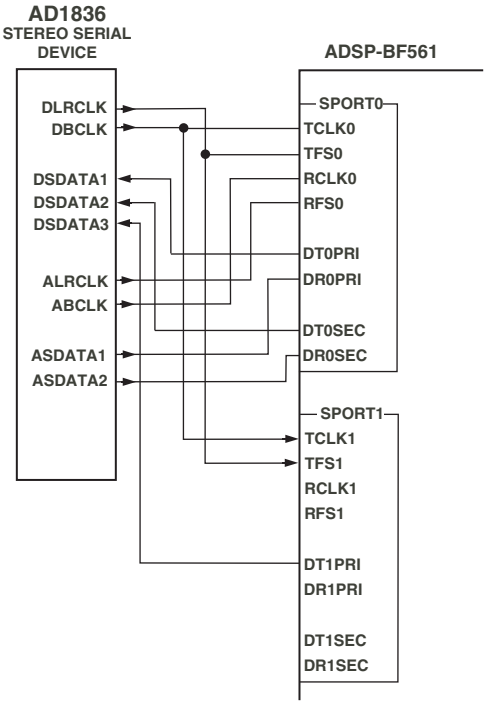

Figure 12-3. Stereo Serial Connection

SPORT Operation

This section describes general SPORT operation, illustrating the most common use of a SPORT. Since the SPORT functionality is configurable, this description represents just one of many possible configurations.

Writing to a SPORT's `SPORTx_TX` register readies the SPORT for transmission. The `TFS` signal initiates the transmission of serial data. Once

transmission has begun, each value written to the SPORTx_TX register is transferred through the FIFO to the internal Transmit Shift register. The bits are then sent, beginning with either the MSB or the LSB as specified in the SPORTx_TCR1 register. Each bit is shifted out on the driving edge of TSCLKx. The driving edge of TSCLKx can be configured to be rising or falling. The SPORT generates the transmit interrupt or requests a DMA transfer as long as there is space in the TX FIFO.

As a SPORT receives bits, they accumulate in an internal receive register. When a complete word has been received, it is written to the SPORT FIFO register and the receive interrupt for that SPORT is generated or a DMA transfer is initiated. Interrupts are generated differently if DMA block transfers are performed. For information about DMA, see Chapter 9, “Direct Memory Access”.

SPORT Disable

The SPORTs are automatically disabled by a processor hardware or software reset. A SPORT can also be disabled directly by clearing the SPORT’s transmit or receive enable bits (TSPEN in the SPORTx_TCR1 register and RSPEN in the SPORTx_RCR1 register, respectively). Each method has a different effect on the SPORT.

A processor reset disables the SPORTs by clearing the SPORTx_TCR1, SPORTx_TCR2, SPORTx_RCR1, and SPORTx_RCR2 registers (including the TSPEN and RSPEN enable bits) and the SPORTx_TCLKDIV, SPORTx_RCLKDIV, SPORTx_TFSDIVx, and SPORTx_RFSDIVx Clock and Frame Sync Divisor registers. Any ongoing operations are aborted.

Clearing the TSPEN and RSPEN enable bits disables the SPORTs and aborts any ongoing operations. Status bits are also cleared. Configuration bits remain unaffected and can be read by the software in order to be altered or

overwritten. To disable the SPORT output clock, set the SPORT to be disabled.

Note that disabling a SPORT via TSPEN/RSPEN may shorten any currently active pulses on the TFSx/RFSx and TSCLKx/RSCLKx pins, if these signals are configured to be generated internally.

The SPORTs are ready to start transmitting or receiving data no later than three serial clock cycles after they are enabled in the SPORTx_TCR1 or SPORTx_RCR1 register. No serial clock cycles are lost from this point on. The first internal frame sync will occur one frame sync delay after the SPORTs are ready. External frame syncs can occur as soon as the SPORT is ready.

When disabling the SPORT from multichannel operation, first disable TXEN and then disable RXEN. Note both TXEN and RXEN must be disabled before re-enabling. Disabling only TX or RX is not allowed.

Setting SPORT Modes

SPORT configuration is accomplished by setting bit and field values in configuration registers. Each SPORT must be configured prior to being enabled. Once the SPORT is enabled, further writes to the SPORT Configuration registers are disabled (except for SPORTx_RCLKDIV, SPORTx_TCLKDIV, and Multichannel Mode Channel Select registers). To change values in all other SPORT Configuration registers, disable the SPORT by clearing TSPEN in SPORTx_TCR1 and/or RSPEN in SPORTx_RCR1.

Each SPORT has its own set of control registers and data buffers. These registers are described in detail in the following sections. All control and status bits in the SPORT registers are active high unless otherwise noted.

Register Writes and Effective Latency

When the SPORT is disabled (TSPEN and RSPEN cleared), SPORT register writes are internally completed at the end of the SCLK cycle in which they occurred, and the register reads back the newly-written value on the next cycle.

When the SPORT is enabled to transmit (TSPEN set) or receive (RSPEN set), corresponding SPORT Configuration register writes are disabled (except for SPORTx_RCLKDIV, SPORTx_TCLKDIV, and Multichannel Mode Channel Select registers). The SPORTx_TX register writes are always enabled; SPORTx_RX, SPORTx_CHNL, and SPORTx_STAT are read-only registers.

After a write to a SPORT register, while the SPORT is disabled, any changes to the control and mode bits generally take effect when the SPORT is re-enabled.

Most configuration registers can only be changed while the SPORT is disabled (TSPEN/RSPEN = 0). Changes take effect after the SPORT is re-enabled. The only exceptions to this rule are the TCLKDIV/RCLKDIV registers and Multichannel Select registers.

SPORTx_TCR1 and SPORTx_TCR2 Registers

The main control registers for the transmit portion of each SPORT are the Transmit Configuration registers, SPORTx_TCR1 and SPORTx_TCR2 (Figure 12-4 and Figure 12-5).

A SPORT is enabled for transmit if Bit 0 (TSPEN) of the Transmit Configuration 1 register is set to 1. This bit is cleared during either a hard reset or a soft reset, disabling all SPORT transmission.

When the SPORT is enabled to transmit (TSPEN set), corresponding SPORT Configuration register writes are not allowed except for SPORTx_TCLKDIV and Multichannel Mode Channel Select registers. Writes to disallowed registers have no effect. While the SPORT is enabled, SPORTx_TCR1 is not written except for bit 0 (TSPEN). For example:

```
write (SPORTx_TCR1, 0x0001) ;   /* SPORT TX Enabled */
write (SPORTx_TCR1, 0xFF01) ;   /* ignored, no effect */
write(SPORTx_TCR1,0xFFF0);   /*SPORTdisabled,SPORTx_TCR1
still equal to 0x0000 */
```

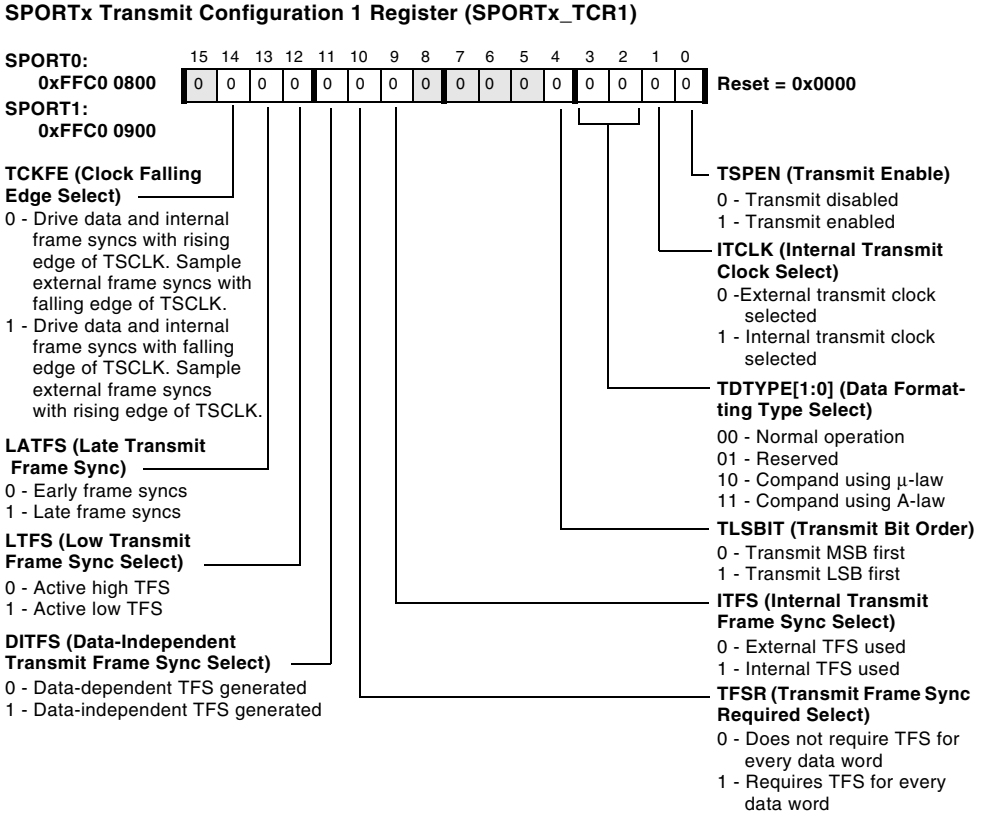

Figure 12-4. SPORTx Transmit Configuration 1 Register

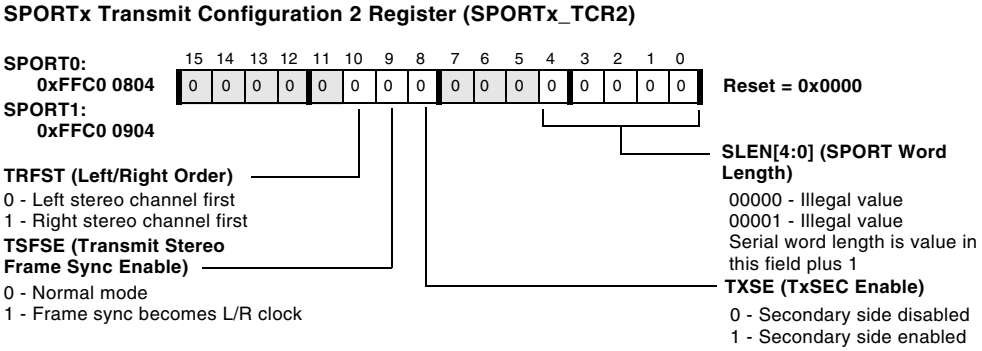

Figure 12-5. SPORTx Transmit Configuration 2 Register

Additional information for the SPORTx_TCR1 and SPORTx_TCR2 Transmit Configuration register bits includes:

- **Transmit Enable** (TSPEN). This bit selects whether the SPORT is enabled to transmit (if set) or disabled (if cleared).

  Setting TSPEN causes an immediate assertion of a SPORT TX interrupt, indicating that the TX data register is empty and needs to be filled. This is normally desirable because it allows centralization of the transmit data write code in the TX interrupt service routine (ISR). For this reason, the code should initialize the ISR and be ready to service TX interrupts before setting TSPEN.

  Similarly, if DMA transfers will be used, DMA control should be configured correctly before setting TSPEN. Set all DMA control registers before setting TSPEN.

  Clearing TSPEN causes the SPORT to stop driving data, TSCLK, and frame sync pins; it also shuts down the internal SPORT circuitry. In low power applications, battery life can be extended by clearing TSPEN whenever the SPORT is not in use.

All SPORT control registers should be programmed before TSPEN is set. Typical SPORT initialization code first writes all control registers, including DMA control if applicable. The last step in the code is to write SPORTx_TCR1 with all of the necessary bits, including TSPEN.

- **Internal Transmit Clock Select**. (ITCLK). This bit selects the internal transmit clock (if set) or the external transmit clock on the TSCLK pin (if cleared). The TCLKDIV MMR value is not used when an external clock is selected.
- **Data Formatting Type Select**. The two TDTYPE bits specify data formats used for single and multichannel operation.
- **Bit Order Select**. (TLSBIT). The TLSBIT bit selects the bit order of the data words transmitted over the SPORT.
- **Serial Word Length Select**. (SLEN). The serial word length (the number of bits in each word transmitted over the SPORTs) is calculated by adding 1 to the value of the SLEN field:

```
Serial Word Length = SLEN + 1;
```

  The SLEN field can be set to a value of 2 to 31; 0 and 1 are illegal values for this field. Three common settings for the SLEN field are 15, to transmit a full 16-bit word; 7, to transmit an 8-bit byte; and 23, to transmit a 24-bit word. The processor can load 16- or 32-bit values into the transmit buffer via DMA or an MMR write instruction; the SLEN field tells the SPORT how many of those bits to shift out of the register over the serial link. The serial port transfers bits [SLEN:0] from the transmit buffer.

The frame sync signal is controlled by the SPORTx_TFSDIV and SPORTx_RFSDIV registers, not by SLEN. To produce a frame sync pulse on each byte or word transmitted, the proper frame sync

divider must be programmed into the Frame Sync Divider register; setting SLEN to 7 does not produce a frame sync pulse on each byte transmitted.

- **Internal Transmit Frame Sync Select**. (ITFS). This bit selects whether the SPORT uses an internal TFS (if set) or an external TFS (if cleared).
- **Transmit Frame Sync Required Select**. (TFSR). This bit selects whether the SPORT requires (if set) or does not require (if cleared) a Transmit Frame Sync for every data word.

The TFSR bit is normally set during SPORT configuration. A frame sync pulse is used to mark the beginning of each word or data packet, and most systems need a frame sync to function properly.

- **Data-Independent Transmit Frame Sync Select**. (DITFS). This bit selects whether the SPORT generates a data-independent TFS (sync at selected interval) or a data-dependent TFS (sync when data is present in SPORTx_TX) for the case of internal frame sync select (ITFS = 1). The DITFS bit is ignored when external frame syncs are selected.

  The frame sync pulse marks the beginning of the data word. If DITFS is set, the frame sync pulse is issued on time, whether the SPORTx_TX register has been loaded or not; if DITFS is cleared, the frame sync pulse is only generated if the SPORTx_TX data register has been loaded. If the receiver demands regular frame sync pulses, DITFS should be set, and the processor should keep loading the SPORTx_TX register on time. If the receiver can tolerate occasional late frame sync pulses, DITFS should be cleared to prevent the SPORT from transmitting old data twice or transmitting garbled data if the processor is late in loading the SPORTx_TX register.
- **Low Transmit Frame Sync Select**. (LTFS). This bit selects an active low TFS (if set) or active high TFS (if cleared).

- **Late Transmit Frame Sync**. (LATFS). This bit configures late frame syncs (if set) or early frame syncs (if cleared).
- **Clock Drive/Sample Edge Select**. (TCKFE). This bit selects which edge of the TCLKx signal the SPORT uses for driving data, for driving internally generated frame syncs, and for sampling externally generated frame syncs. If set, data and internally generated frame syncs are driven on the falling edge, and externally generated frame syncs are sampled on the rising edge. If cleared, data and internally generated frame syncs are driven on the rising edge, and externally generated frame syncs are sampled on the falling edge.
- **TxSec Enable**. (TXSE). This bit enables the transmit secondary side of the serial port (if set).
- **Stereo Serial Enable**. (TSFSE). This bit enables the Stereo Serial operating mode of the serial port (if set). By default this bit is cleared, enabling normal clocking and frame sync.
- **Left/Right Order**. (TRFST). If this bit is set, the right channel is transmitted first in Stereo Serial operating mode. By default this bit is cleared, and the left channel is transmitted first.

SPORTx_RCR1 and SPORTx_RCR2 Registers

The main control registers for the receive portion of each SPORT are the Receive Configuration registers, SPORTx_RCR1 and SPORTx_RCR2 (Figure 12-6 and Figure 12-7).

A SPORT is enabled for receive if Bit 0 (RSPEN) of the Receive Configuration 1 register is set to 1. This bit is cleared during either a hard reset or a soft reset, disabling all SPORT reception.

When the SPORT is enabled to receive (RSPEN set), corresponding SPORT Configuration register writes are not allowed except for SPORTx_RCLKDIV and Multichannel Mode Channel Select registers. Writes to disallowed registers have no effect. While the SPORT is enabled, SPORTx_RCR1 is not written except for bit 0 (RSPEN). For example:

```
write (SPORTx_RCR1, 0x0001) ;  /* SPORT RX Enabled */
write (SPORTx_RCR1, 0xFF01) ;  /* ignored, no effect */
write (SPORTx_RCR1, 0xFFF0) ;  /* SPORT disabled, SPORTx_RCR1
still equal to 0x0000 */
```

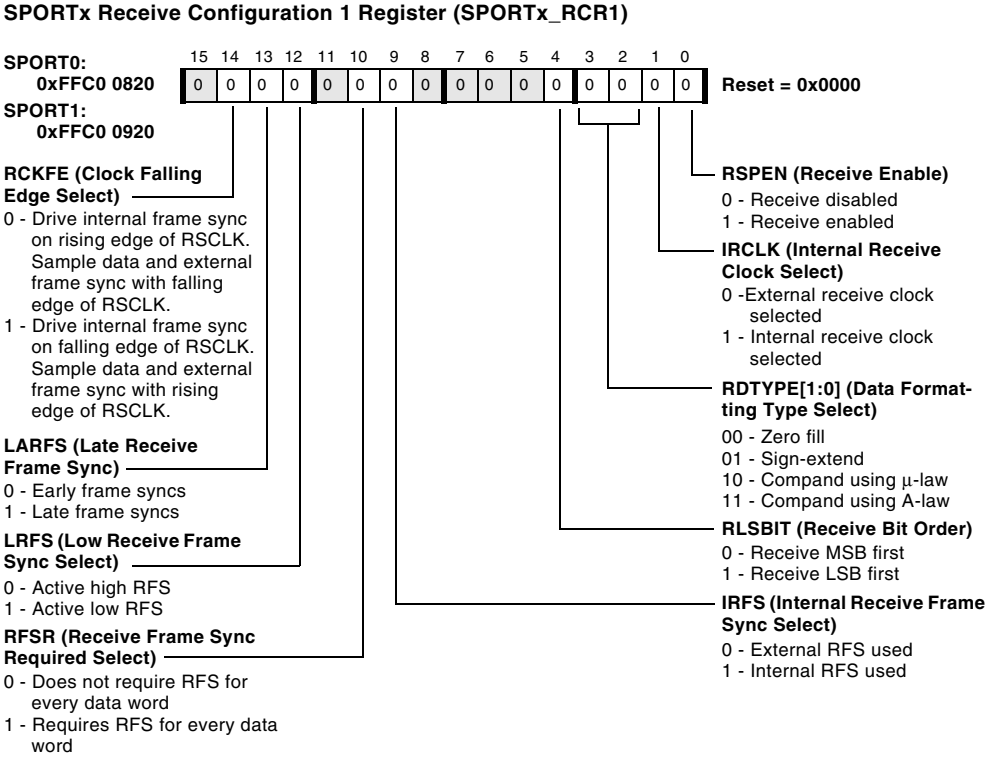

Figure 12-6. SPORTx Receive Configuration 1 Register

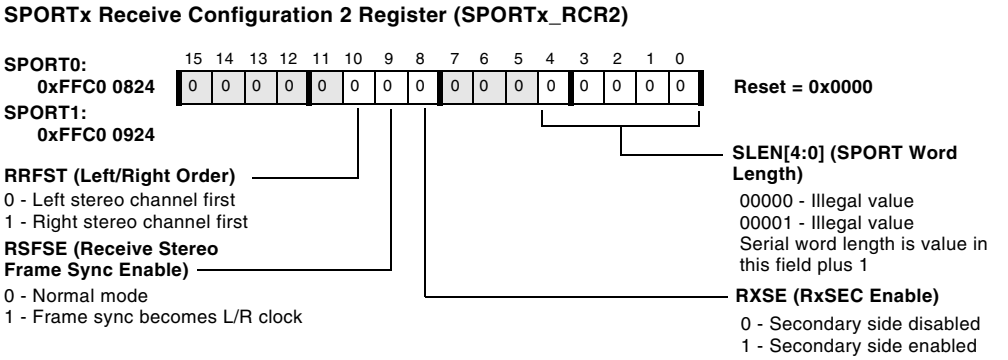

Figure 12-7. SPORTx Receive Configuration 2 Register

Additional information for the SPORTx_RCR1 and SPORTxRCR2 Receive Configuration register bits:

- **Receive Enable**. (RSPEN). This bit selects whether the SPORT is enabled to receive (if set) or disabled (if cleared). Setting the RSPEN bit turns on the SPORT and causes it to sample data from the data receive pins as well as the receive bit clock and Receive Frame Sync pins if so programmed.

  Setting RSPEN enables the SPORTx receiver, which can generate a SPORTx RX interrupt. For this reason, the code should initialize the ISR and the DMA control registers, and should be ready to service RX interrupts before setting RSPEN. Setting RSPEN also generates DMA requests if DMA is enabled and data is received. Set all DMA control registers before setting RSPEN.

  Clearing RSPEN causes the SPORT to stop receiving data; it also shuts down the internal SPORT receive circuitry. In low power applications, battery life can be extended by clearing RSPEN whenever the SPORT is not in use.

All SPORT control registers should be programmed before RSPEN is set. Typical SPORT initialization code first writes all control registers, including DMA control if applicable. The last step in the code is to write SPORTx_RCR1 with all of the necessary bits, including RSPEN.

- **Internal Receive Clock Select**. (IRCLK). This bit selects the internal receive clock (if set) or external receive clock (if cleared). The RCLKDIV MMR value is not used when an external clock is selected.
- **Data Formatting Type Select**. (RDTYPE). The two RDTYPE bits specify one of four data formats used for single and multichannel operation.
- **Bit Order Select**. (RLSBIT). The RLSBIT bit selects the bit order of the data words received over the SPORTs.
- **Serial Word Length Select**. (SLEN). The serial word length (the number of bits in each word received over the SPORTs) is calculated by adding 1 to the value of the SLEN field. The SLEN field can be set to a value of 2 to 31; 0 and 1 are illegal values for this field.

The frame sync signal is controlled by the SPORTx_TFSDIV and SPORTx_RFSDIV registers, not by SLEN. To produce a frame sync pulse on each byte or word transmitted, the proper frame sync divider must be programmed into the Frame Sync Divider register; setting SLEN to 7 does not produce a frame sync pulse on each byte transmitted.

- **Internal Receive Frame Sync Select**. (IRFS). This bit selects whether the SPORT uses an internal RFS (if set) or an external RFS (if cleared).
- **Receive Frame Sync Required Select**. (RFSR). This bit selects whether the SPORT requires (if set) or does not require (if cleared) a Receive Frame Sync for every data word.

- **Low Receive Frame Sync Select**. (LRFS). This bit selects an active low RFS (if set) or active high RFS (if cleared).
- **Late Receive Frame Sync**. (LARFS). This bit configures late frame syncs (if set) or early frame syncs (if cleared).
- **Clock Drive/Sample Edge Select**. (RCKFE). This bit selects which edge of the RSCLK clock signal the SPORT uses for sampling data, for sampling externally generated frame syncs, and for driving internally generated frame syncs. If set, internally generated frame syncs are driven on the falling edge, and data and externally generated frame syncs are sampled on the rising edge. If cleared, internally generated frame syncs are driven on the rising edge, and data and externally generated frame syncs are sampled on the falling edge.
- **RxSec Enable**. (RXSE). This bit enables the receive secondary side of the serial port (if set).
- **Stereo Serial Enable**. (RSFSE). This bit enables the Stereo Serial operating mode of the serial port (if set). By default this bit is cleared, enabling normal clocking and frame sync.
- **Left/Right Order**. (RRFST). If this bit is set, the right channel is received first in Stereo Serial operating mode. By default this bit is cleared, and the left channel is received first.

Data Word Formats

The format of the data words transferred over the SPORTs is configured by the combination of transmit SLEN and receive SLEN; RDTYPE; TDTYPE; RLSBIT; and TLSBIT bits of the SPORTx_TCR1, SPORTx_TCR2, SPORTx_RCR1, and SPORTx_RCR2 registers.

SPORTx_TX Register

The SPORTx Transmit Data register (SPORTx_TX) is a write-only register (see Figure 12-8). Reads produce a Peripheral Access Bus (PAB) error (see Figure 12-9).

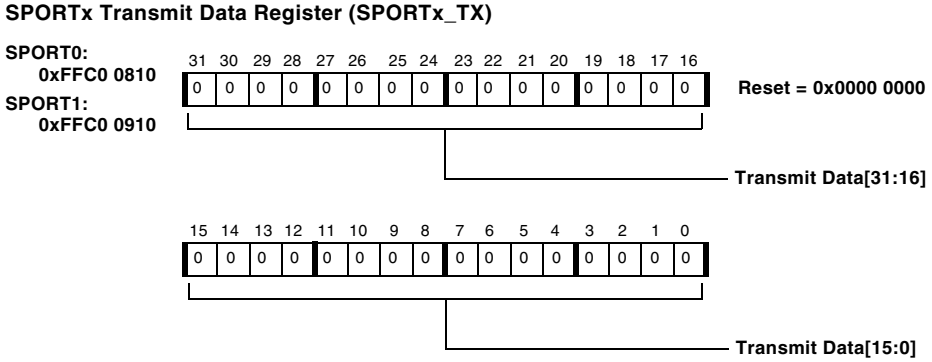

Figure 12-8. SPORTx Transmit Data Register

Writes to this register cause writes into the transmitter FIFO. The 16-bit wide FIFO is 8 deep for word length <= 16 and 4 deep for word length > 16. The FIFO is common to both primary and secondary data and stores data for both. Data ordering in the FIFO is shown in the Figure 12-9.

It is important to keep the interleaving of primary and secondary data in the FIFO as shown. This means that PAB/DMA writes to the FIFO must follow an order of primary first, and then secondary, if secondary is enabled. DAB/PAB writes must match their size to the data word length. For word length up to and including 16 bits, use a 16-bit write. Use a 32-bit write for word length greater than 16 bits.

When transmit is enabled, data from the FIFO is assembled in the TX Hold register based on TXSE and SLEN, and then shifted into the primary and secondary shift registers. From here, the data is shifted out serially on the DTPRI and DTSEC pins.

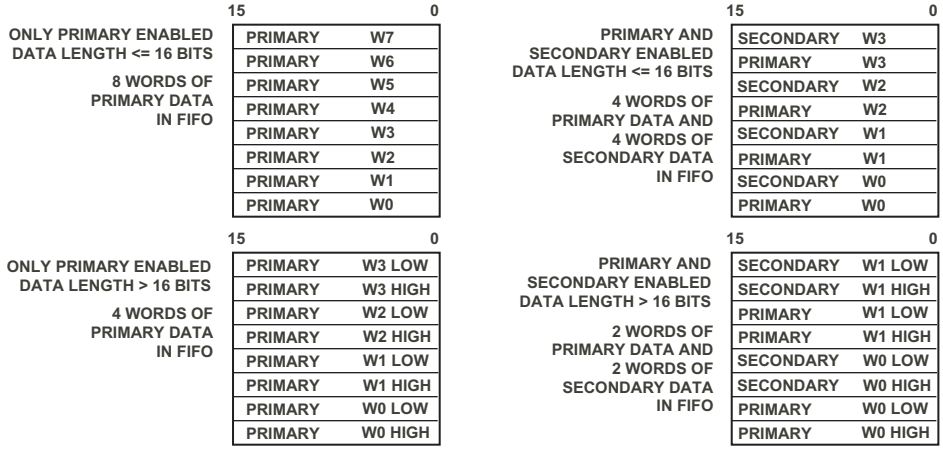

Figure 12-9. SPORT Transmit FIFO Data Ordering

The SPORT TX interrupt is asserted when `TSPEN = 1` and the TX FIFO has room for additional words. This interrupt does not occur if SPORT DMA is enabled. For DMA operation, see Chapter 9, “Direct Memory Access”.

The Transmit Underflow Status bit (`TUVF`) is set in the SPORT Status register when a Transmit Frame Sync occurs and no new data has been loaded into the serial shift register. In Multichannel Mode (MCM), `TUVF` is set whenever the serial shift register is not loaded, and transmission begins on the current enabled channel. The `TUVF` status bit is a sticky write-1-to-clear (W1C) bit and is also cleared by disabling the serial port (writing `TXEN = 0`).

If software causes the core processor to attempt a write to a full TX FIFO with a `SPORTx_TX` write, the new data is lost and no overwrites occur to data in the FIFO. The `TOVF` status bit is set and a SPORT error interrupt is asserted. The `TOVF` bit is a sticky bit; it is only cleared by disabling the SPORT TX. To find out whether the core processor can access the `SPORTx_TX` register without causing this type of error, read the register’s

status first. The TXF bit in the SPORT Status register is 0 if space is available for another word in the FIFO.

The TXF and TOVF status bits in the SPORTx Status register are updated upon writes from the core processor, even when the SPORT is disabled.

SPORTx_RX Register

The SPORTx Receive Data register (SPORTx_RX) is a read-only register (see Figure 12-10). Writes produce a PAB error. The same location is read for both primary and secondary data. Reading from this register space causes reading of the receive FIFO. This 16-bit FIFO is 8 deep for receive word length <= 16 and 4 deep for length > 16 bits. The FIFO is shared by both primary and secondary receive data. The order for reading using PAB/DMA reads is important since data is stored in differently depending on the setting of the SLEN and RXSE configuration bits.

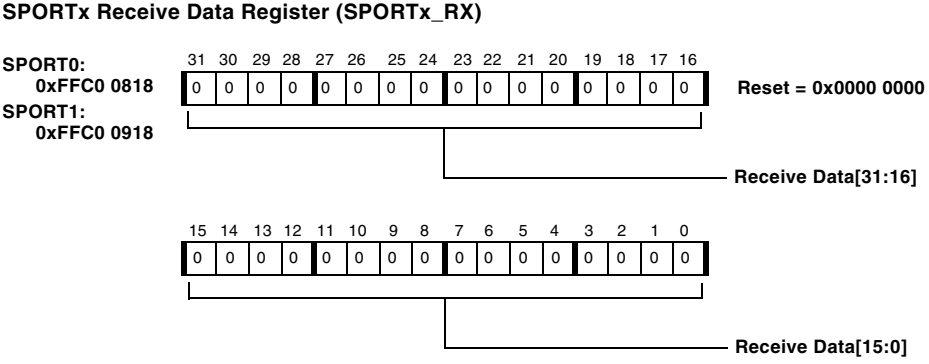

Figure 12-10. SPORTx Receive Data Register

Data storage and data ordering in the FIFO is shown in Figure 12-11.

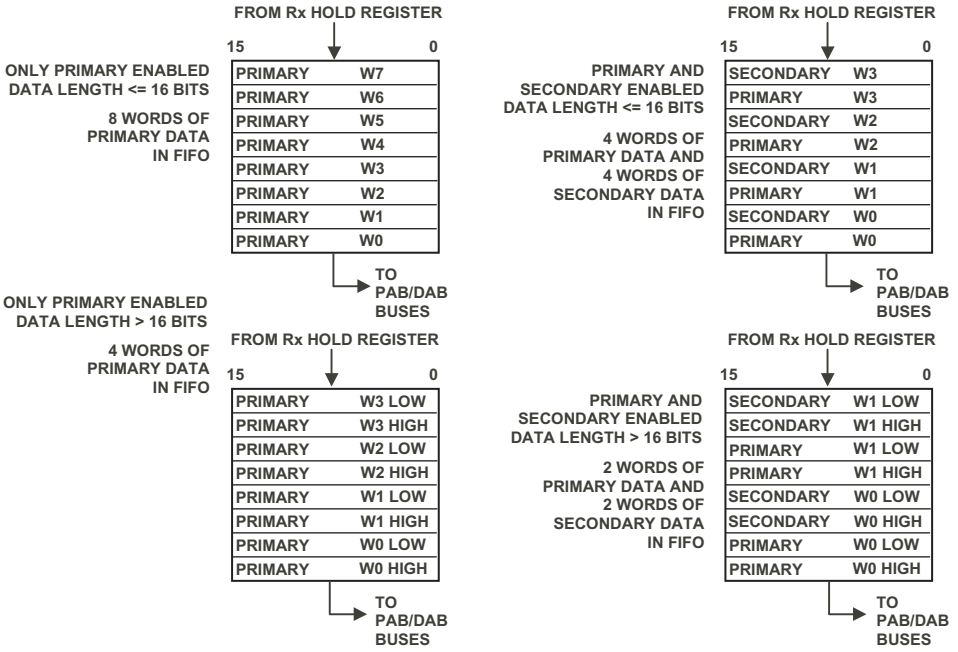

Figure 12-11. SPORT Receive FIFO Data Ordering

When reading from the FIFO for both primary and secondary data, read primary first, followed by secondary. DAB/PAB reads must match their size to the data word length. For word length up to and including 16 bits, use a 16-bit read. Use a 32-bit read for word length greater than 16 bits.

When receiving is enabled, data from the DRPRI pin is loaded into the RX primary shift register, while data from the DRSEC pin is loaded into the RX secondary shift register. At transfer completion of a word, data is shifted into the RX Hold registers for primary and secondary data, respectively. Data from the Hold registers is moved into the FIFO based on RXSE and SLEN.

The SPORT RX interrupt is generated when RSPEN = 1 and the RX FIFO has received words in it. When the core processor has read all the words in the FIFO, the RX interrupt is cleared. The SPORT RX interrupt is set only if SPORT RX DMA is disabled; otherwise, the FIFO is read by DMA reads.

If the program causes the core processor to attempt a read from an empty RX FIFO, old data is read, the RUVF flag is set in the SPORTx_STAT register, and the SPORT error interrupt is asserted. The RUVF bit is a sticky bit and is cleared only when the SPORT is disabled. To determine if the core can access the RX registers without causing this error, first read the RX FIFO status (RXNE in the SPORTx Status register). The RUVF status bit is updated even when the SPORT is disabled.

The ROVF status bit is set in the SPORTx_STAT register when a new word is assembled in the RX Shift register and the RX Hold register has not moved the data to the FIFO. The previously written word in the Hold register is overwritten. The ROVF bit is a sticky bit; it is only cleared by disabling the SPORT RX.

SPORTx_STAT Register

The SPORT Status register (SPORTx_STAT) is used to determine if the access to a SPORT RX or TX FIFO can be made by determining their full or empty status (see Figure 12-12).

The TXF bit in the SPORT Status register indicates whether there is room in the TX FIFO. The RXNE status bit indicates whether there are words in the RX FIFO. The TXHRE bit indicates if the TX Hold register is empty.

The Transmit Underflow Status bit (TUVF) is set whenever the TFS signal occurs (from either an external or internal source) while the TX Shift register is empty. The internally generated TFS may be suppressed whenever SPORTx_TX is empty by clearing the DITFS control bit in the SPORT Configuration register. The TUVF status bit is a sticky write-1-to-clear

(W1C) bit and is also cleared by disabling the serial port (writing TXEN = 0).

For continuous transmission (TFSR = 0), TUVF is set at the end of a transmitted word if no new word is available in the TX Hold register.

The TOVF bit is set when a word is written to the TX FIFO when it is full. It is a sticky W1C bit and is also cleared by writing TXEN = 0. Both TXF and TOVF are updated even when the SPORT is disabled.

When the SPORT RX Hold register is full, and a new receive word is received in the Shift register, the Receive Overflow Status bit (ROVF) is set in the SPORT Status register. It is a sticky W1C bit and is also cleared by disabling the serial port (writing RXEN = 0).

The RUVF bit is set when a read is attempted from the RX FIFO and it is empty. It is a sticky W1C bit and is also cleared by writing RXEN = 0. The RUVF bit is updated even when the SPORT is disabled.

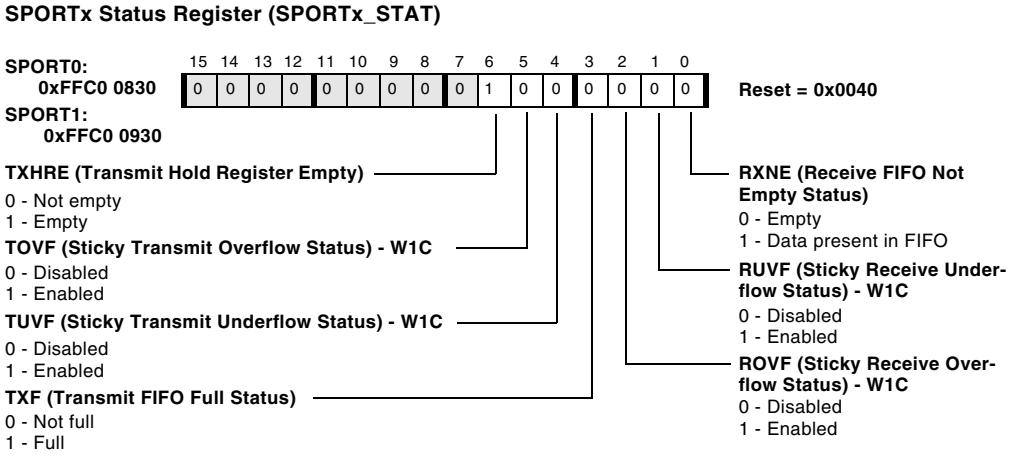

Figure 12-12. SPORTx Status Register

SPORT RX, TX, and Error Interrupts

The SPORT RX interrupt is asserted when RSPEN is enabled and any words are present in the RX FIFO. If RX DMA is enabled, the SPORT RX interrupt is turned off and DMA services the RX FIFO.

The SPORT TX interrupt is asserted when TSPEN is enabled and the TX FIFO has room for words. If TX DMA is enabled, the SPORT TX interrupt is turned off and DMA services the TX FIFO.

The SPORT Error interrupt is asserted when any of the sticky status bits (ROVF, RUVF, TOVF, TUVF) are set. The ROVF and RUVF bits are cleared by writing 0 to RSPEN. The TOVF and TUVF bits are cleared by writing 0 to TSPEN.

PAB Errors

The SPORT generates a PAB error for illegal register read or write operations. Examples include:

- Reading a write-only register (for example, SPORT_TX)
- Writing a read-only register (for example, SPORT_RX)
- Writing or reading a register with the wrong size (for example, 32-bit read of a 16-bit register)
- Accessing reserved register locations

SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers

The frequency of an internally generated clock is a function of the system clock frequency (as seen at the SCLK pin) and the value of the 16-bit serial clock divide modulus registers (the SPORTx Transmit Serial Clock

Divider register, SPORTx_TCLKDIV, and the SPORTx Receive Serial Clock Divider register, SPORTx_RCLKDIV). See Figure 12-13 and Figure 12-14.

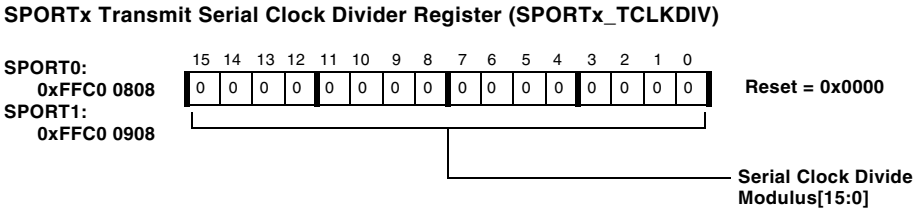

Figure 12-13. SPORTx Transmit Serial Clock Divider Register

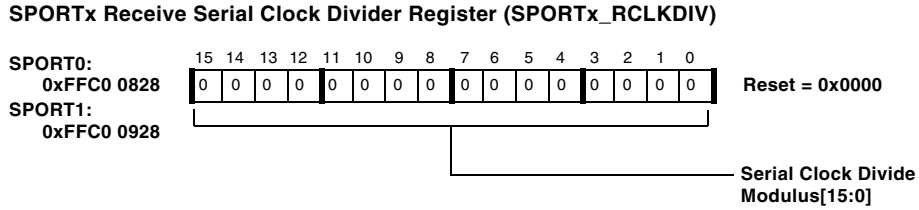

Figure 12-14. SPORTx Receive Serial Clock Divider Register

SPORTx_TFSDIV and SPORTx_RFSDIV Register

The 16-bit SPORTx Transmit Frame Sync Divider register (SPORTx_TFSDIV) and the SPORTx Receive Frame Sync Divider register (SPORTx_RFSDIV) specify how many transmit or receive clock cycles are counted before generating a TFS or RFS pulse when the frame sync is internally generated (Figure 12-15 and Figure 12-16). In this way, a frame sync can be used to initiate periodic transfers. The counting of serial clock cycles applies to either internally or externally generated serial clocks.

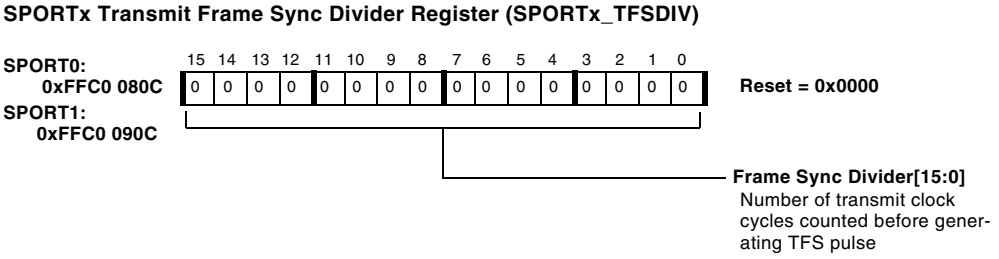

Figure 12-15. SPORTx Transmit Frame Sync Divider Register

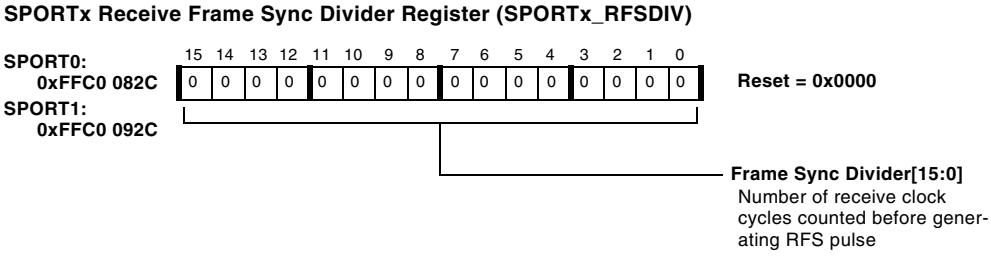

Figure 12-16. SPORTx Receive Frame Sync Divider Register

Clock and Frame Sync Frequencies

The maximum serial clock frequency (for either an internal source or an external source) is SCLK/2. The frequency of an internally generated clock is a function of the system clock frequency (SCLK) and the value of the 16-bit serial clock divide modulus registers, SPORTx_TCLKDIV and SPORTx_RCLKDIV.

```
SPORTx_TCLK frequency =
(SCLK frequency)/(2 × (SPORTx_TCLKDIV + 1))

SPORTx_RCLK frequency =
(SCLK frequency)/(2 × (SPORTx_RCLKDIV + 1))
```

If the value of SPORTx_TCLKDIV or SPORTx_RCLKDIV is changed while the internal serial clock is enabled, the change in TSCLK or RSCLK frequency takes effect at the start of the drive edge of TSCLK or RSCLK that follows the next leading edge of TFS or RFS.

When an internal frame sync is selected (ITFS = 1 in the SPORTx_TCR1 register or IRFS = 1 in the SPORTx_RCR1 register) and frame syncs are not required, the first frame sync does not update the clock divider if the value in SPORTx_TCLKDIV or SPORTx_RCLKDIV has changed. The second frame sync will cause the update.

The SPORTx_TFSDIV and SPORTx_RFSDIV registers specify the number of transmit or receive clock cycles that are counted before generating a TFS or RFS pulse (when the frame sync is internally generated). This enables a frame sync to initiate periodic transfers. The counting of serial clock cycles applies to either internally or externally generated serial clocks.

The formula for the number of cycles between frame sync pulses is:

```
# of transmit serial clocks between frame sync assertions =
TFSDIV + 1
# of receive serial clocks between frame sync assertions =
RFSDIV + 1
```

Use the following equations to determine the correct value of TFSDIV or RFSDIV, given the serial clock frequency and desired frame sync frequency:

```
SPORTxTFS frequency = (TSCLKx frequency)/(SPORTx_TFSDIV + 1)

SPORTxRFS frequency = (RSCLKx frequency)/(SPORTx_RFSDIV + 1)
```

The frame sync would thus be continuously active (for transmit if TFSDIV = 0 or for receive if RFSDIV = 0). However, the value of TFSDIV (or RFSDIV) should not be less than the serial word length minus 1 (the value of the SLEN field in SPORTx_TCR2 or SPORTx_RCR2). A smaller value could cause an external device to abort the current operation or have other unpredictable results. If a SPORT is not being used, the TFSDIV

(or RFSDIV) divisor can be used as a counter for dividing an external clock or for generating a periodic pulse or periodic interrupt. The SPORT must be enabled for this mode of operation to work.

Maximum Clock Rate Restrictions

Externally generated late Transmit Frame Syncs also experience a delay from arrival to data output, and this can limit the maximum serial clock speed. See the *ADSP-BF561 Embedded Processor Data Sheet* for exact timing specifications.

Frame Sync and Clock Example

The following code fragment is a brief example of setting up the clocks and frame sync.

```
r0 = 0x00FF;
p0.l = SPORT0_RFSDIV & 0xFFFF;
p0.h = (SPORT0_RFSDIV >> 16) & 0xFFFF;
w[p0] = r0.l; ssync;
p0.l = SPORT0_TFSDIV & 0xFFFF;
w[p0] = r0.l; ssync;
```

Word Length

Each SPORT channel (transmit and receive) independently handles word lengths of 3 to 32 bits. The data is right-justified in the SPORT data registers if it is fewer than 32 bits long, residing in the LSB positions. The value of the serial word length (SLEN) field in the SPORTx_TCR2 and SPORTx_RCR2 registers of each SPORT determines the word length according to this formula:

```
Serial Word Length = SLEN + 1
```

The SLEN value should not be set to 0 or 1; values from 2 to 31 are allowed. Continuous operation (when the last bit of the current word is immediately followed by the first bit of the next word) is restricted to word sizes of 4 or longer (so SLEN ≥ 3).

Bit Order

Bit order determines whether the serial word is transmitted MSB first or LSB first. Bit order is selected by the RLSBIT and TLSBIT bits in the SPORTx_RCR1 and SPORTx_TCR1 registers. When RLSBIT (or TLSBIT) = 0, serial words are received (or transmitted) MSB first. When RLSBIT (or TLSBIT) = 1, serial words are received (or transmitted) LSB first.

Data Type

The TDTYPE field of the SPORTx_TCR1 register and the RDTYPE field of the SPORTx_RCR1 register specify one of four data formats for both single and multichannel operation. See Table 12-2.

Table 12-2. TDTYPE, RDTYPE, and Data Formatting

TDTYPE or RDTYPE	SPORTx_TCR1 Data Formatting	SPORTx_RCR1 Data Formatting
00	Normal operation	Zero fill
01	Reserved	Sign extend
10	Compand using μ-law	Compand using μ-law
11	Compand using A-law	Compand using A-law

These formats are applied to serial data words loaded into the SPORTx_RX and SPORTx_TX buffers. SPORTx_TX data words are not actually zero filled or sign extended, because only the significant bits are transmitted.

Companding

Companding (a contraction of COMpressing and exPANDing) is the process of logarithmically encoding and decoding data to minimize the number of bits that must be sent. The SPORTs support the two most widely used companding algorithms, µ-law and A-law. The processor compands data according to the CCITT G.711 specification. The type of companding can be selected independently for each SPORT.

When companding is enabled, valid data in the SPORTx_RX register is the right-justified, expanded value of the eight LSBs received and sign extended to 16 bits. A write to SPORTx_TX causes the 16-bit value to be compressed to eight LSBs (sign extended to the width of the transmit word) and written to the internal transmit register. Although the companding standards support only 13-bit (A-law) or 14-bit (µ-law) maximum word lengths, up to 16-bit word lengths can be used. If the magnitude of the word value is greater than the maximum allowed, the value is automatically compressed to the maximum positive or negative value.

Lengths greater than 16 bits are not supported for companding operation.

Clock Signal Options

Each SPORT has a transmit clock signal (TSCLK) and a receive clock signal (RSCLK). The clock signals are configured by the TCKFE and RCKFE bits of the SPORTx_TCR1 and SPORTx_RCR1 registers. Serial clock frequency is configured in the SPORTx_TCLKDIV and SPORTx_RCLKDIV registers.

The receive clock pin may be tied to the transmit clock if a single clock is desired for both receive and transmit.

Both transmit and receive clocks can be independently generated internally or input from an external source. The ITCLK bit of the SPORTx_TCR1

Configuration register and the IRCLK bit in the SPORTx_RCR1 Configuration register determines the clock source.

When IRCLK or ITCLK = 1, the clock signal is generated internally by the core, and the TSCLK or RSCLK pin is an output. The clock frequency is determined by the value of the serial clock divisor in the SPORTx_RCLKDIV register.

When IRCLK or ITCLK = 0, the clock signal is accepted as an input on the TSCLK or RSCLK pins, and the serial clock divisors in the SPORTx_TCLKDIV/SPORTx_RCLKDIV registers are ignored. The externally generated serial clocks do not need to be synchronous with the core system clock or with each other. The core system clock must have a higher frequency than RSCLK and TSCLK.

Frame Sync Options

Framing signals indicate the beginning of each serial word transfer. The framing signals for each SPORT are TFS (Transmit Frame Sync) and RFS (Receive Frame Sync). A variety of framing options are available; these options are configured in the SPORT Configuration registers (SPORTx_TCR1, SPORTx_TCR2, SPORTx_RCR1 and SPORTx_RCR2). The TFS and RFS signals of a SPORT are independent and are separately configured in the control registers.

Framed Versus Unframed

The use of multiple frame sync signals is optional in SPORT communications. The TFSR (Transmit Frame Sync Required Select) and RFSR (Receive Frame Sync Required Select) control bits determine whether frame sync signals are required. These bits are located in the SPORTx_TCR1 and SPORTx_RCR1 registers.

When TFSR = 1 or RFSR = 1, a frame sync signal is required for every data word. To allow continuous transmitting by the SPORT, each new data word must be loaded into the SPORTx_TX Hold register before the previous word is shifted out and transmitted.

When TFSR = 0 or RFSR = 0, the corresponding frame sync signal is not required. A single frame sync is needed to initiate communications but is ignored after the first bit is transferred. Data words are then transferred continuously, unframed.

With frame syncs not required, interrupt or DMA requests may not be serviced frequently enough to guarantee continuous unframed data flow. Monitor status bits or check for a SPORT Error interrupt to detect underflow or overflow of data.

Figure 12-17 illustrates framed serial transfers, which have these characteristics:

- TFSR and RFSR bits in the SPORTx_TCR1 and SPORTx_RCR1 registers determine framed or unframed mode.
- Framed mode requires a framing signal for every word. Unframed mode ignores a framing signal after the first word.
- Unframed mode is appropriate for continuous reception.
- Active low or active high frame syncs are selected with the LTFS and LRFS bits of the SPORTx_TCR1 and SPORTx_RCR1 registers.

See “Timing Examples” on page 12-62 for more timing examples.

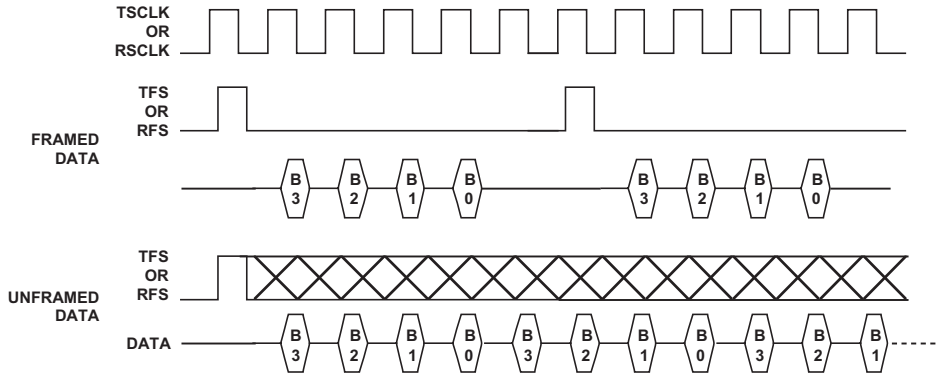

Figure 12-17. Framed Versus Unframed Data

Internal Versus External Frame Syncs

Both Transmit and Receive Frame Syncs can be independently generated internally or can be input from an external source. The ITFS and IRFS bits of the SPORTx_TCR1 and SPORTx_RCR1 registers determine the frame sync source.

When ITFS = 1 or IRFS = 1, the corresponding frame sync signal is generated internally by the SPORT, and the TFS pin or RFS pin is an output. The frequency of the frame sync signal is determined by the value of the frame sync divisor in the SPORTx_TFSDIV or SPORTx_RFSDIV register.

When ITFS = 0 or IRFS = 0, the corresponding frame sync signal is accepted as an input on the TFS pin or RFS pin, and the frame sync divisors in the SPORTx_TFSDIV/SPORTx_RFSDIV registers are ignored.

All of the frame sync options are available whether the signal is generated internally or externally.

Active Low Versus Active High Frame Syncs

Frame sync signals may be either active high or active low (in other words, inverted). The LTFS and LRFS bits of the SPORTx_TCR1 and SPORTx_RCR1 registers determine frame sync logic levels:

- When LTFS = 0 or LRFS = 0, the corresponding frame sync signal is active high.
- When LTFS = 1 or LRFS = 1, the corresponding frame sync signal is active low.

Active high frame syncs are the default. The LTFS and LRFS bits are initialized to 0 after a processor reset.

Sampling Edge for Data and Frame Syncs

Data and frame syncs can be sampled on either the rising or falling edges of the SPORT clock signals. The TCKFE and RCKFE bits of the SPORTx_TCR1 and SPORTx_RCR1 registers select the driving and sampling edges of the serial data and frame syncs.

For the SPORT transmitter, setting TCKFE = 1 in the SPORTx_TCR1 register selects the falling edge of TSCLKx to drive data and internally generated frame syncs and selects the rising edge of TSCLKx to sample externally generated frame syncs. Setting TCKFE = 0 selects the rising edge of TSCLKx to drive data and internally generated frame syncs and selects the falling edge of TSCLKx to sample externally generated frame syncs.

For the SPORT receiver, setting RCKFE = 1 in the SPORTx_RCR1 register selects the falling edge of RSCLKx to drive internally generated frame syncs and selects the rising edge of RSCLKx to sample data and externally generated frame syncs. Setting RCKFE = 0 selects the rising edge of RSCLKx to

drive internally generated frame syncs and selects the falling edge of RSCLKx to sample data and externally generated frame syncs.

Note externally generated data and frame sync signals should change state on the opposite edge than that selected for sampling. For example, for an externally generated frame sync to be sampled on the rising edge of the clock (TCKFE = 1 in the SPORTx_TCR1 register), the frame sync must be driven on the falling edge of the clock.

The transmit and receive functions of two SPORTs connected together should always select the same value for TCKFE in the transmitter and RCKFE in the receiver, so that the transmitter drives the data on one edge and the receiver samples the data on the opposite edge.

In Figure 12-18, TCKFE = RCKFE = 0 and transmit and receive are connected together to share the same clock and frame syncs.

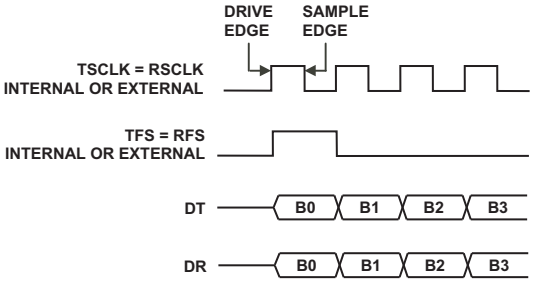

Figure 12-18. Example of TCKFE = RCKFE = 0, Transmit and Receive Connected

In Figure 12-19, TCKFE = RCKFE = 1 and transmit and receive are connected together to share the same clock and frame syncs.

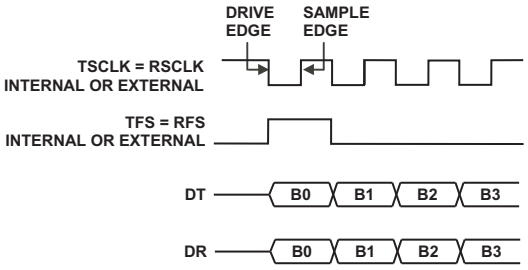

Figure 12-19. Example of TCKFE = RCKFE = 1, Transmit and Receive Connected

Early Versus Late Frame Syncs (Normal Versus Alternate Timing)

Frame sync signals can occur during the first bit of each data word (late) or during the serial clock cycle immediately preceding the first bit (early). The LATFS and LARFS bits of the SPORTx_TCR1 and SPORTx_RCR1 registers configure this option.

When LATFS = 0 or LARFS = 0, early frame syncs are configured; this is the normal mode of operation. In this mode, the first bit of the transmit data word is available and the first bit of the receive data word is sampled in the serial clock cycle after the frame sync is asserted, and the frame sync is not checked again until the entire word has been transmitted or received. In multichannel operation, this corresponds to the case when multichannel frame delay is 1.

If data transmission is continuous in early framing mode (in other words, the last bit of each word is immediately followed by the first bit of the next word), then the frame sync signal occurs during the last bit of each word.

Internally generated frame syncs are asserted for one clock cycle in early framing mode. Continuous operation is restricted to word sizes of 4 or longer (SLEN ≥ 3).

When LATFS = 1 or LARFS = 1, late frame syncs are configured; this is the alternate mode of operation. In this mode, the first bit of the transmit data word is available and the first bit of the receive data word is sampled in the same serial clock cycle that the frame sync is asserted. In multichannel operation, this is the case when frame delay is 0. Receive data bits are sampled by serial clock edges, but the frame sync signal is only checked during the first bit of each word. Internally generated frame syncs remain asserted for the entire length of the data word in late framing mode. Externally generated frame syncs are only checked during the first bit.

Figure 12-20 illustrates the two modes of frame signal timing. In summary:

- For the LATFS or LARFS bits of the SPORTx_TCR1 or SPORTx_RCR1 registers: LATFS = 0 or LARFS = 0 for early frame syncs, LATFS = 1 or LARFS = 1 for late frame syncs.
- For early framing, the frame sync precedes data by one cycle. For late framing, the frame sync is checked on the first bit only.
- Data is transmitted MSB first (TLSBIT = 0 or RLSBIT = 0) or LSB first (TLSBIT = 1 or RLSBIT = 1).
- Frame sync and clock are generated internally or externally.

See "Timing Examples" on page 12-62 for more examples.

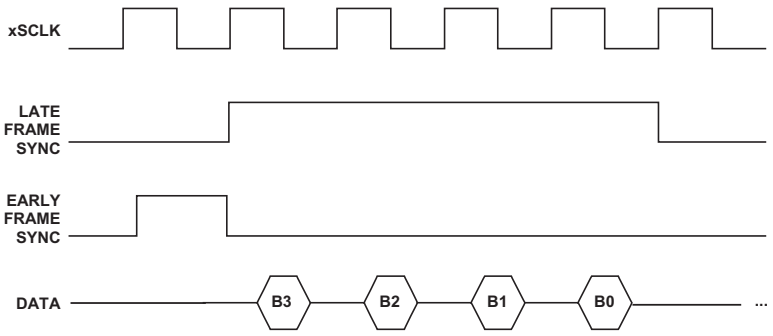

Figure 12-20. Normal Versus Alternate Framing

Data Independent Transmit Frame Sync

Normally the internally generated Transmit Frame Sync signal (TFS) is output only when the SPORTx_TX buffer has data ready to transmit. The Data-Independent Transmit Frame Sync Select bit (DITFS) allows the continuous generation of the TFS signal, with or without new data. The DITFS bit of the SPORTx_TCR1 register configures this option.

When DITFS = 0, the internally generated TFS is only output when a new data word has been loaded into the SPORTx_TX buffer. The next TFS is generated once data is loaded into SPORTx_TX. This mode of operation allows data to be transmitted only when it is available.

When DITFS = 1, the internally generated TFS is output at its programmed interval regardless of whether new data is available in the SPORTx_TX buffer. Whatever data is present in SPORTx_TX is transmitted again with each assertion of TFS. The TUVF (Transmit Underflow Status) bit in the SPORTx_STAT register is set when this occurs and old data is retransmitted. The TUVF status bit is also set if the SPORTx_TX buffer does not have new

data when an externally generated TFS occurs. Note that in this mode of operation, data is transmitted only at specified times.

If the internally generated TFS is used, a single write to the SPORTx_TX data register is required to start the transfer.

Moving Data Between SPORTs and Memory

Transmit and receive data can be transferred between the SPORTs and on-chip memory in one of two ways: with single word transfers or with DMA block transfers.

If no SPORT DMA channel is enabled, the SPORT generates an interrupt every time it has received a data word or needs a data word to transmit. SPORT DMA provides a mechanism for receiving or transmitting an entire block or multiple blocks of serial data before the interrupt is generated. The SPORT's DMA controller handles the DMA transfer, allowing the processor core to continue running until the entire block of data is transmitted or received. Interrupt service routines (ISRs) can then operate on the block of data rather than on single words, significantly reducing overhead.

For information about DMA, see Chapter 9, "Direct Memory Access".

Stereo Serial Operation

Several Stereo Serial modes can be supported by the SPORT, including the popular I^2S format. To use these modes, set bits in the SPORT_RCR2 or SPORT_TCR2 registers. Setting RSFSE or TSFSE in SPORT_RCR2 or SPORT_TCR2 changes the operation of the frame sync pin to a left/right clock as required for I^2S and left-justified stereo serial data. Setting this bit enables

the SPORT to generate or accept the special LRCLK-style frame sync. All other SPORT control bits remain in effect and should be set appropriately. Figure 12-21 on page 12-45 and Figure 12-22 on page 12-46 show timing diagrams for Stereo Serial mode operation.

Table 12-3 shows several modes that can be configured using bits in SPORTx_TCR1 and SPORTx_RCR1. The table shows bits for the receive side of the SPORT, but corresponding bits are available for configuring the transmit portion of the SPORT. A control field which may be either set or cleared depending on the user's needs, without changing the standard, is indicated by an "X."

Table 12-3. Stereo Serial Settings

Bit Field	Stereo Audio Serial Scheme		
	I^2S	Left-Justified	DSP Mode
RSFSE	1	1	0
RRFST	0	0	0
LARFS	0	1	0
LRFS	0	1	0
RFSR	1	1	1
RCKFE	1	0	0
SLEN	2 – 31	2 – 31	2 – 31
RLSBIT	0	0	0
RFSDIV (If internal FS is selected.)	2 – Max	2 – Max	2 – Max
RXSE (Secondary Enable is available for RX and TX.)	X	X	X

Note most bits shown as a 0 or 1 may be changed depending on the user's preference, creating many other "almost standard" modes of stereo serial operation. These modes may be of use in interfacing to codecs with

slightly non-standard interfaces. The settings shown in Table 12-3 provide glueless interfaces to many popular codecs.

Note RFSDIV or TFSDIV must still be greater than or equal to SLEN. For I^2S operation, RFSDIV or TFSDIV is usually 1/64 of the serial clock rate. With RSFSE set, the formulas to calculate frame sync period and frequency (discussed in "Clock and Frame Sync Frequencies" on page 12-29) still apply, but now refer to one half the period and twice the frequency. For instance, setting RFSDIV or TFSDIV = 31 produces an LRCLK that transitions every 32 serial clock cycles and has a period of 64 serial clock cycles.

The LRFS bit determines the polarity of the RFS or TFS frame sync pin for the channel that is considered a "right" channel. Thus, setting LRFS = 0 (meaning that it is an active high signal) indicates that the frame sync is high for the "right" channel, thus implying that it is low for the "left" channel. This is the default setting.

The RRFST and TRFST bits determine whether the first word received or transmitted is a left or a right channel. If the bit is set, the first word received or transmitted is a right channel. The default is to receive or transmit the left channel word first.

The secondary DRxSEC and DTxSEC pins are useful extensions of the serial port which pair well with Stereo Serial mode. Multiple I^2S streams of data can be transmitted or received using a single SPORT. Note the primary and secondary pins are synchronous, as they share clock and LRCLK (Frame Sync) pins. The transmit and receive sides of the SPORT need not be synchronous, but may share a single clock in some designs. See Figure 12-3 on page 12-8, which shows multiple stereo serial connections being made between the processor and an AD1836 codec.

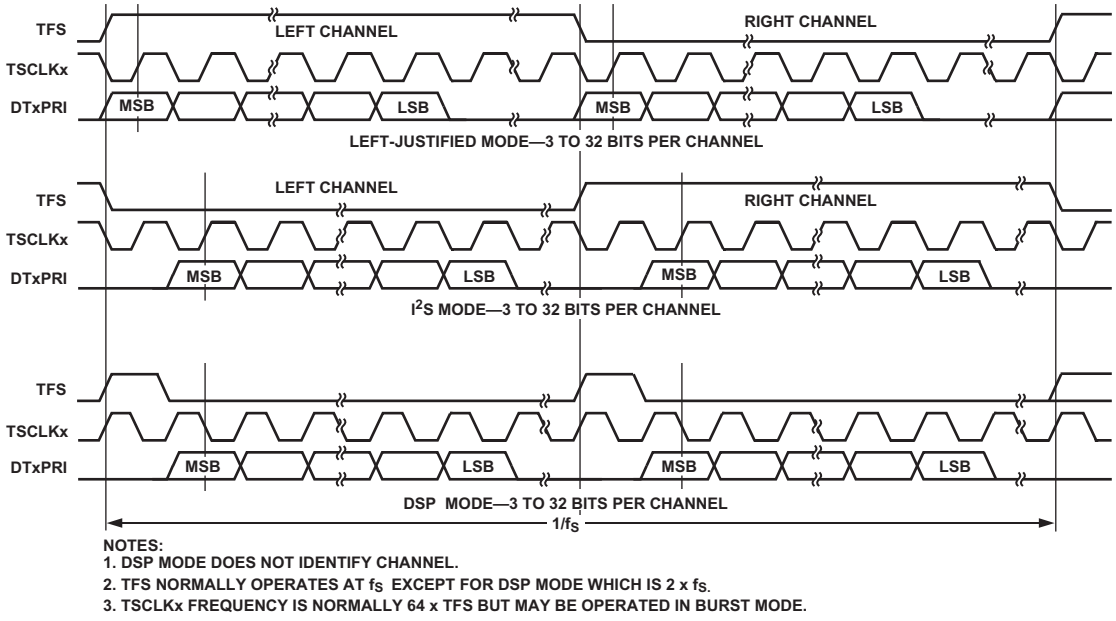

Figure 12-21. SPORT Stereo Serial Modes, Transmit

The Blackfin processor's SPORTs are designed such that in I^{2}S master mode, the LRCLK signal is held at the last driven logic level. The LRCLK signal does not transistion (provide an edge) after the final data word is driven out. While transmitting a fixed number of words to an I^{2}S receiver that expects an LRCLK edge to receive the incoming data word, the SPORT should send a dummy word after transmitting the fixed number of words. The transmission of this dummy word toggles LRCLK, generating an edge. Transmission of the dummy word is not required when the I^{2}S receiver is a Blackfin processor SPORT.

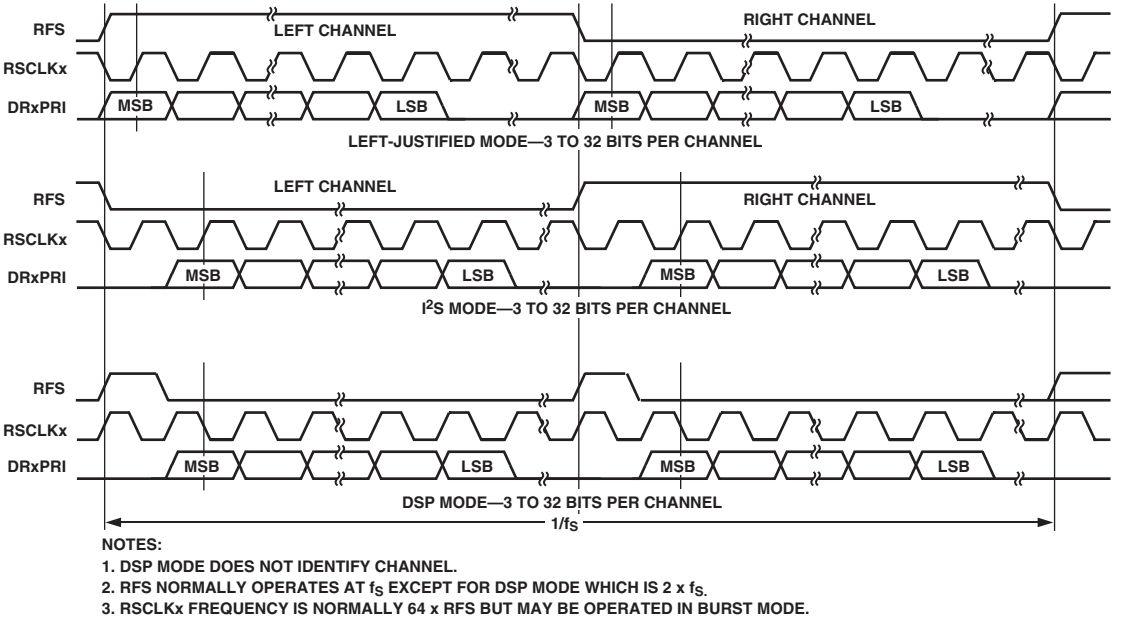

Figure 12-22. SPORT Stereo Serial Modes, Receive

Multichannel Operation

The SPORTs offer a Multichannel mode of operation which allows the SPORT to communicate in a Time-Division-Multiplexed (TDM) serial system. In multichannel communications, each data word of the serial bit-stream occupies a separate channel. Each word belongs to the next consecutive channel so that, for example, a 24-word block of data contains one word for each of 24 channels.

The SPORT can automatically select words for particular channels while ignoring the others. Up to 128 channels are available for transmitting or receiving; each SPORT can receive and transmit data selectively from any of the 128 channels. These 128 channels can be any 128 out of the 1024 total channels. RX and TX must use the same 128-channel region to

selectively enable channels. The SPORT can do any of the following on each channel:

- Transmit data
- Receive data
- Transmit and receive data
- Do nothing

Data companding and DMA transfers can also be used in multichannel mode.

The DTPRI pin is always driven (not three-stated) if the SPORT is enabled (TSPEN = 1 in the SPORTx_TCR1 register), unless it is in Multichannel mode and an inactive time slot occurs. The DTSEC pin is always driven (not three-stated) if the SPORT is enabled and the secondary transmit is enabled (TXSE = 1 in the SPORTx_TCR2 register), unless the SPORT is in Multichannel mode and an inactive time slot occurs.

In Multichannel mode, RSCLK can either be provided externally or generated internally by the SPORT, and it is used for both transmit and receive functions. Leave TSCLK disconnected if the SPORT is used only in multichannel mode. If RSCLK is externally or internally provided, it will be internally distributed to both the receiver and transmitter circuitry.

The SPORT Multichannel Transmit Select register and the SPORT Multichannel Receive Select register must be programmed before enabling SPORTx_TX or SPORTx_RX operation for Multichannel Mode. This is especially important in “DMA data unpacked mode,” since SPORT FIFO operation begins immediately after RSPEN and TSPEN are set, enabling both RX and TX. The MCMEN bit (in SPORTx_MCMC2) must be enabled prior to enabling SPORTx_TX or SPORTx_RX operation. When disabling the SPORT from

multichannel operation, first disable TXEN and then disable RXEN. Note both TXEN and RXEN must be disabled before re-enabling. Disabling only TX or RX is not allowed.

Figure 12-23 shows example timing for a multichannel transfer that has these characteristics:

- Use TDM method where serial data is sent or received on different channels sharing the same serial bus
- Can independently select transmit and receive channels
- RFS signals start of frame
- TFS is used as "Transmit Data Valid" for external logic, true only during transmit channels
- Receive on channels 0 and 2, transmit on channels 1 and 2
- Multichannel frame delay is set to 1

See "Timing Examples" on page 12-62 for more examples.

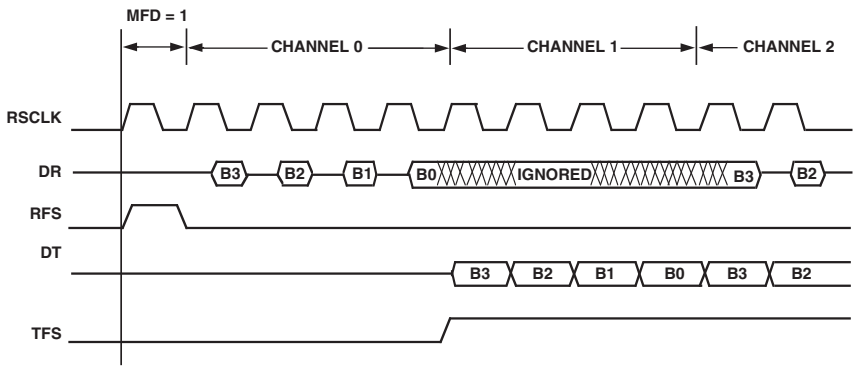

Figure 12-23. Multichannel Operation

SPORTx_MCMCn Registers

There are two SPORTx Multichannel Configuration registers (SPORTx_MCMCn) for each SPORT. The SPORTx_MCMCn registers are used to configure the multichannel operation of the SPORT. The two control registers are shown in Figure 12-24 and Figure 12-25.

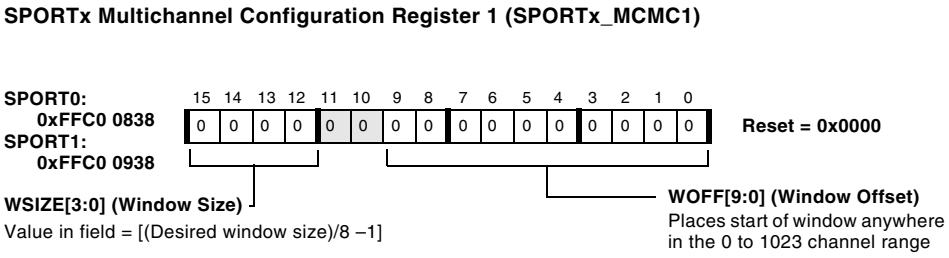

Figure 12-24. SPORTx Multichannel Configuration Register 1

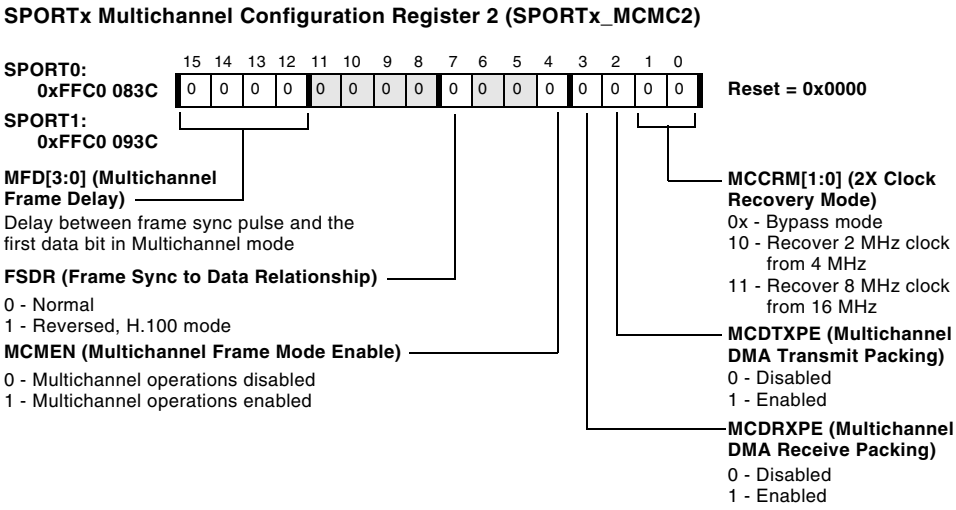

Figure 12-25. SPORTx Multichannel Configuration Register 2

Multichannel Enable

Setting the MCMEN bit in the SPORTx_MCM2 register enables Multichannel mode. When MCMEN = 1, multichannel operation is enabled; when MCMEN = 0, all multichannel operations are disabled.

Setting the MCMEN bit enables multichannel operation for *both* the receive and transmit sides of the SPORT. Therefore, if a receiving SPORT is in Multichannel mode, the transmitting SPORT must also be in Multichannel mode.

When in multichannel mode, do not enable the stereo serial frame sync modes or the late frame sync feature, as these features are incompatible with multichannel mode.

Table 12-4 shows the dependencies of bits in the SPORT Configuration register when the SPORT is in Multichannel Mode.

Table 12-4. Multichannel Mode Configuration

SPORTx_RCR1 or SPORTx_RCR2	SPORTx_TCR1 or SPORTx_TCR2	Notes
RSPEN	TSPEN	Set or clear both
IRCLK	-	Independent
-	ITCLK	Ignored
RDTYPE	TDTYPE	Independent
RLSBIT	TLSBIT	Independent
IRFS	-	Independent
-	ITFS	Ignored
RFSR	TFSR	Ignored
-	DITFS	Ignored
LRFS	LTFS	Independent
LARFS	LATFS	Both must be 0

Table 12-4. Multichannel Mode Configuration (Cont'd)

SPORTx_RCR1 or SPORTx_RCR2	SPORTx_TCR1 or SPORTx_TCR2	Notes
RCKFE	TCKFE	Set or clear both to same value
SLEN	SLEN	Set or clear both to same value
RXSE	TXSE	Independent
RSFSE	TSFSE	Both must be 0
RRFST	TRFST	Ignored

Frame Syncs in Multichannel Mode

All receiving and transmitting devices in a multichannel system must have the same timing reference. The RFS signal is used for this reference, indicating the start of a block or frame of multichannel data words.

When Multichannel mode is enabled on a SPORT, both the transmitter and the receiver use RFS as a frame sync. This is true whether RFS is generated internally or externally. The RFS signal is used to synchronize the channels and restart each multichannel sequence. Assertion of RFS indicates the beginning of the channel 0 data word.

Since RFS is used by both the SPORTx_TX and SPORTx_RX channels of the SPORT in Multichannel mode configuration, the corresponding bit pairs in SPORTx_RCR1 and SPORTx_TCR1, and in SPORTx_RCR2 and SPORTx_TCR2, should always be programmed identically, with the possible exception of the RXSE and TXSE pair and the RDTYPE and TDTYPE pair. This is true even if SPORTx_RX operation is not enabled.

In Multichannel mode, RFS timing similar to late (alternative) frame mode is entered automatically; the first bit of the transmit data word is available and the first bit of the receive data word is sampled in the same serial clock cycle that the frame sync is asserted, provided that MFD is set to 0.

The TFS signal is used as a transmit data valid signal which is active during transmission of an enabled word. The SPORT's data transmit pin is three-stated when the time slot is not active, and the TFS signal serves as an output-enabled signal for the data transmit pin. The SPORT drives TFS in Multichannel mode whether or not ITFS is cleared. The TFS pin in Multichannel mode still obeys the LTFS bit. If LTFS is set, the transmit data valid signal will be active low—a low signal on the TFS pin indicates an active channel.

Once the initial RFS is received, and a frame transfer has started, all other RFS signals are ignored by the SPORT until the complete frame has been transferred.

If MFD > 0, the RFS may occur during the last channels of a previous frame. This is acceptable, and the frame sync is not ignored as long as the delayed channel 0 starting point falls outside the complete frame.

In Multichannel mode, the RFS signal is used for the block or frame start reference, after which the word transfers are performed continuously with no further RFS signals required. Therefore, internally generated frame syncs are always data independent.

The Multichannel Frame

A multichannel frame contains more than one channel, as specified by the window size and window offset. A complete multichannel frame consists of 1 – 1024 channels, starting with channel 0. The particular channels of the multichannel frame that are selected for the SPORT are a combination of the window offset, the window size, and the Multichannel Select registers. See Figure 12-26.

Multichannel Frame Delay

The 4-bit MFD field in SPORTx_MCMC2 specifies a delay between the frame sync pulse and the first data bit in Multichannel mode. The value of MFD is

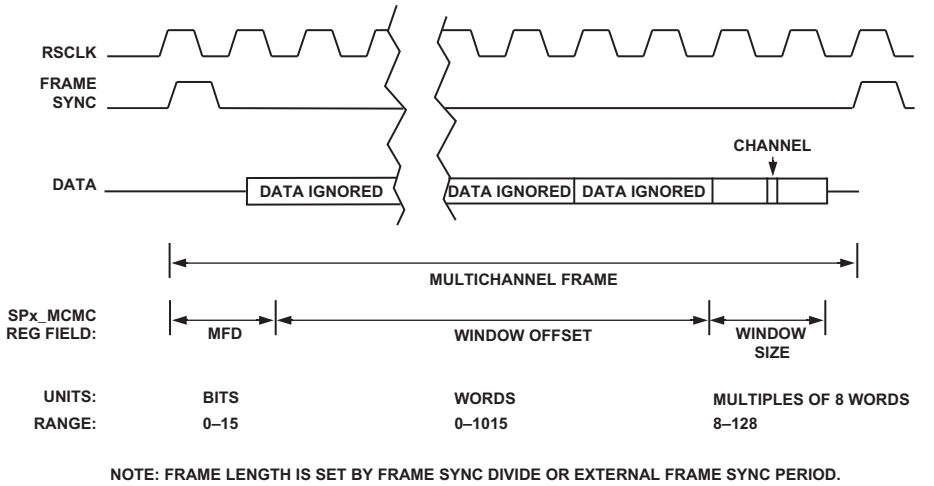

Figure 12-26. Relationships for Multichannel Parameters

the number of serial clock cycles of the delay. Multichannel frame delay allows the processor to work with different types of interface devices.

A value of 0 for MFD causes the frame sync to be concurrent with the first data bit. The maximum value allowed for MFD is 15. A new frame sync may occur before data from the last frame has been received, because blocks of data occur back-to-back.

Window Size

The window size (WSIZE[3:0]) defines the number of channels that can be enabled/disabled by the Multichannel Select registers. This range of words is called the active window. The number of channels can be any value in the range of 0 to 15, corresponding to active window size of 8 to 128, in increments of 8; the default value of 0 corresponds to a minimum active window size of 8 channels. To calculate the active window size from the WSIZE register, use this equation:

```
Number of words in active window = 8 × (WSIZE + 1)
```

Since the DMA buffer size is always fixed, it is possible to define a smaller window size (for example, 32 words), resulting in a smaller DMA buffer size (in this example, 32 words instead of 128 words) to save DMA bandwidth. The window size cannot be changed while the SPORT is enabled.

Multichannel select bits that are enabled but fall outside the window selected are ignored.

Window Offset

The window offset (WOFF[9:0]) specifies where in the 1024-channel range to place the start of the active window. A value of 0 specifies no offset and 896 is the largest value that permits using all 128 channels. As an example, a program could define an active window with a window size of 8 (WSIZE = 0) and an offset of 93 (WOFF = 93). This 8-channel window would reside in the range from 93 to 100. Neither the window offset nor the window size can be changed while the SPORT is enabled.

If the combination of the window size and the window offset would place any portion of the window outside of the range of the channel counter, none of the out-of-range channels in the frame are enabled.

SPORTx_CHNL Register

The 10-bit CHNL field in the SPORTx Current Channel register (SPORTx_CHNL) indicates which channel is currently being serviced during multichannel operation (see Figure 12-27). This field is a read-only status indicator. The CHNL[9:0] field increments by one as each channel is serviced. The counter stops at the upper end of the defined window. The Channel Select register restarts at 0 at each frame sync. As an example, for a window size of 8 and an offset of 148, the counter displays a value between 0 and 156.

Once the window size has completed, the channel counter resets to 0 in preparation for the next frame. Because there are synchronization delays

between RSCLK and the processor clock, the Channel register value is approximate. It is never ahead of the channel being served, but it may lag behind.

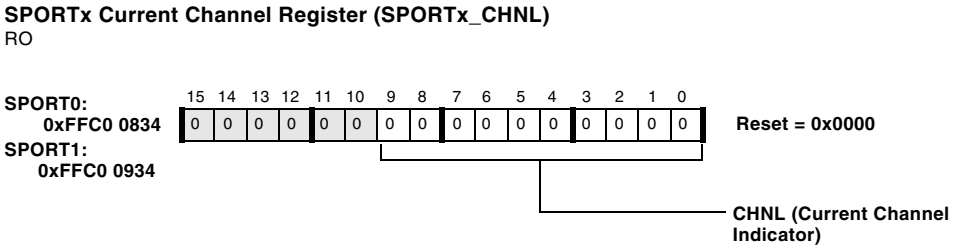

Figure 12-27. SPORTx Current Channel Register

Other Multichannel Fields in SPORTx_MCMC2

The FSDR bit in the SPORTx_MCMC2 register changes the timing relationship between the frame sync and the clock received. This change enables the SPORT to comply with the H.100 protocol.

Normally (When FSDR = 0), the data is transmitted on the same edge that the TFS is generated. For example, a positive edge on TFS causes data to be transmitted on the positive edge of the TSCLK—either the same edge or the following one, depending on when LATFS is set.

When the frame sync/data relationship is used (FSDR = 1), the frame sync is expected to change on the falling edge of the clock and is sampled on the rising edge of the clock. This is true even though data received is sampled on the negative edge of the receive clock.

Channel Selection Register

A channel is a multi-bit word from 3 to 32 bits in length that belongs to one of the TDM channels. Specific channels can be individually enabled or disabled to select which words are received and transmitted during

multichannel communications. Data words from the enabled channels are received or transmitted, while disabled channel words are ignored. Up to 128 contiguous channels may be selected out of 1024 available channels. The SPORTx_MRCSn and SPORTx_MTCSn Multichannel Select registers are used to enable and disable individual channels; the SPORTx_MRCSn registers specify the active receive channels, and the SPORTx_MTCSn registers specify the active transmit channels.

Four registers make up each Multichannel Select register (Figure 12-28). Each of the four registers has 32 bits, corresponding to 32 channels. Setting a bit enables that channel, so the SPORT selects its word from the multiple word block of data (for either receive or transmit).

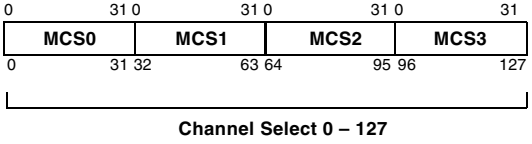

Figure 12-28. Multichannel Select Registers

Channel Select bit 0 always corresponds to the first word of the active window. To determine a channel's absolute position in the frame, add the window offset words to the channel select position. For example, setting bit 7 in MCS2 selects word 71 of the active window to be enabled. Setting bit 2 in MCS1 selects word 34 of the active window, and so on.

Setting a particular bit in the SPORTx_MTCSn register causes the SPORT to transmit the word in that channel's position of the datastream. Clearing the bit in the SPORTx_MTCSn register causes the SPORT's data transmit pin to three-state during the time slot of that channel.

Setting a particular bit in the SPORTx_MRCSn register causes the SPORT to receive the word in that channel's position of the datastream; the received

word is loaded into the SPORTx_RX buffer. Clearing the bit in the SPORTx_MRCSn register causes the SPORT to ignore the data.

Companding may be selected for all channels or for no channels. A-law or μ-law companding is selected with the TDTYPE field in the SPORTx_TCR1 register and the RDTYPE field in the SPORTx_RCR1 register, and applies to all active channels. (See "Companding" on page 12-33 for more information about companding.)

SPORTx_MRCSn Registers

The Multichannel Selection registers are used to enable and disable individual channels. The SPORTx Multichannel Receive Select registers (SPORTx_MRCSn) specify the active receive channels (Figure 12-29).

SPORTx Multichannel Receive Select Registers (SPORTx_MRCSn)

For all bits, 0 - Channel disabled, 1 - Channel enabled, so SPORT selects that word from multiple word block of data.

For memory-mapped addresses, see Table 12-5.

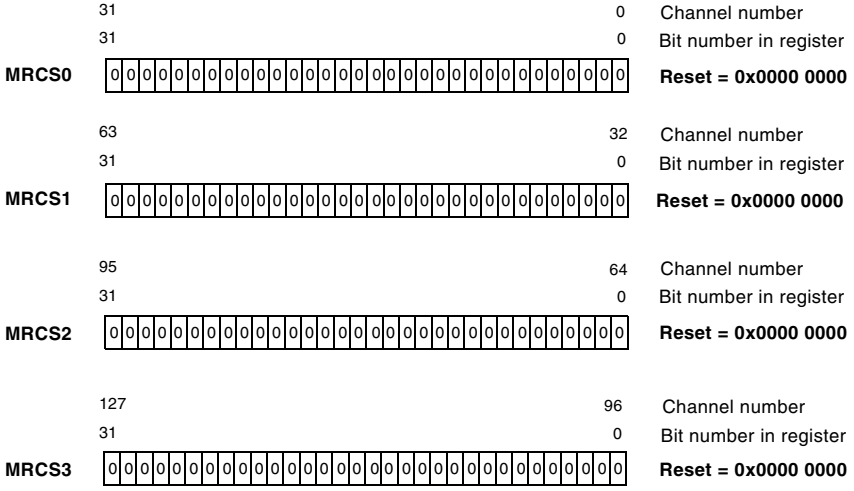

Figure 12-29. SPORTx Multichannel Receive Select Registers

Table 12-5. SPORTx Multichannel Receive Select Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
SPORT0_MRCS0	0xFFC0 0850
SPORT0_MRCS1	0xFFC0 0854
SPORT0_MRCS2	0xFFC0 0858
SPORT0_MRCS3	0xFFC0 085C
SPORT1_MRCS0	0xFFC0 0950
SPORT1_MRCS1	0xFFC0 0954
SPORT1_MRCS2	0xFFC0 0958
SPORT1_MRCS3	0xFFC0 095C

There are four registers, each with 32 bits, corresponding to the 128 channels. Setting a bit enables that channel so that the serial port selects that word for receive from the multiple word block of data. For example, setting bit 0 selects word 0, setting bit 12 selects word 12, and so on.

Setting a particular bit in the SPORTx_MRCSn register causes the serial port to receive the word in that channel's position of the datastream; the received word is loaded into the RX buffer. Clearing the bit in the SPORTx_MRCSn register causes the serial port to ignore the data.

SPORTx_MTCSn Registers

The Multichannel Selection registers are used to enable and disable individual channels. The four SPORTx Multichannel Transmit Select registers (SPORTx_MTCSn) specify the active transmit channels (Figure 12-30). There are four registers, each with 32 bits, corresponding to the 128 channels. Setting a bit enables that channel so that the serial port selects that word for transmit from the multiple word block of data. For example, setting bit 0 selects word 0, setting bit 12 selects word 12, and so on.

Setting a particular bit in a SPORTx_MTCSn register causes the serial port to transmit the word in that channel's position of the datastream. Clearing the bit in the SPORTx_MTCSn register causes the serial port's data transmit pin to three-state during the time slot of that channel.

SPORTx Multichannel Transmit Select Registers (SPORTx_MTCSn)

For all bits, 0 - Channel disabled, 1 - Channel enabled, so SPORT selects that word from multiple word block of data.

For memory-mapped addresses, see Table 12-6.

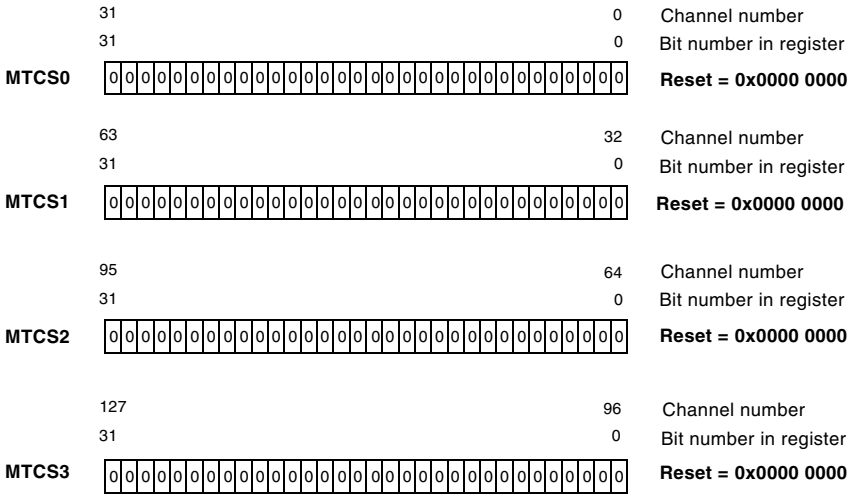

Figure 12-30. SPORTx Multichannel Transmit Select Registers

Table 12-6. SPORTx Multichannel Transmit Select Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
SPORT0_MTCS0	0xFFC0 0840
SPORT0_MTCS1	0xFFC0 0844
SPORT0_MTCS2	0xFFC0 0848

Table 12-6. SPORTx Multichannel Transmit Select Register Memory-Mapped Addresses (Cont'd)

Register Name	Memory-Mapped Address
SPORT0_MTCS3	0xFFC0 084C
SPORT1_MTCS0	0xFFC0 0940
SPORT1_MTCS1	0xFFC0 0944
SPORT1_MTCS2	0xFFC0 0948
SPORT1_MTCS3	0xFFC0 094C

Multichannel DMA Data Packing

Multichannel DMA data packing and unpacking are specified with the MCDTXPE and MCDRXPE bits in the SPORTx_MCMC2 Multichannel Configuration register.

If the bits are set, indicating that data is packed, the SPORT expects the data contained by the DMA buffer corresponds only to the enabled SPORT channels. For example, if an MCM frame contains 10 enabled channels, the SPORT expects the DMA buffer to contain 10 consecutive words for each frame. It is not possible to change the total number of enabled channels without changing the DMA buffer size, and reconfiguring is not allowed while the SPORT is enabled.

If the bits are cleared (the default, indicating that data is not packed), the SPORT expects the DMA buffer to have a word for each of the channels in the active window, whether enabled or not, so the DMA buffer size must be equal to the size of the window. For example, if channels 1 and 10 are enabled, and the window size is 16, the DMA buffer size would have to be 16 words. The data to be transmitted or received would be placed at addresses 1 and 10 of the buffer, and the rest of the words in the DMA buffer would be ignored. This mode allows changing the number of enabled channels while the SPORT is enabled, with some caution. First read the Channel register to make sure that the active window is not being

serviced. If the channel count is 0, then the Multichannel Select registers can be updated.

Support for H.100 Standard Protocol

The processor supports the H.100 standard protocol. The following SPORT parameters must be set to support this standard.

- Set for external frame sync. Frame sync generated by external bus master.
- TFSR/RFSR set (frame syncs required)
- LTFS/LRFS set (active low frame syncs)
- Set for external clock
- MCMEN set (Multichannel mode selected)
- MFD = 0 (no frame delay between frame sync and first data bit)
- SLEN = 7 (8-bit words)
- FSDR = 1 (set for H.100 configuration, enabling half-clock-cycle early frame sync)

2X Clock Recovery Control

The SPORTs can recover the data rate clock from a provided 2X input clock. This enables the implementation of H.100 compatibility modes for MVIP-90 (2 Mbps data) and HMVIP (8 Mbps data), by recovering 2 MHz from 4 MHz or 8 MHz from the 16 MHz incoming clock with the proper phase relationship. A 2-bit mode signal (MCCRM[1:0] in the SPORTx_MCMC2 register) chooses the applicable clock mode, which includes a non-divide or bypass mode for normal operation. A value of MCCRM = 00 chooses non-divide or bypass mode (H.100-compatible), MCCRM = 10

chooses MVIP-90 clock divide (extract 2 MHz from 4 MHz), and `MCCRM = 11` chooses HMVIP clock divide (extract 8 MHz from 16 MHz).

SPORT Pin/Line Terminations

The processor has very fast drivers on all output pins, including the SPORTs. If connections on the data, clock, or frame sync lines are longer than six inches, consider using a series termination for strip lines on point-to-point connections. This may be necessary even when using low speed serial clocks, because of the edge rates.

Timing Examples

Several timing examples are included within the text of this chapter (in the sections "Framed Versus Unframed" on page 12-34, "Early Versus Late Frame Syncs (Normal Versus Alternate Timing)" on page 12-39, and "Frame Syncs in Multichannel Mode" on page 12-51). This section contains additional examples to illustrate other possible combinations of the framing options.

These timing examples show the relationships between the signals but are not scaled to show the actual timing parameters of the processor. Consult the *ADSP-BF561 Embedded Processor Data Sheet* for actual timing parameters and values.

These examples assume a word length of four bits (`SLEN = 3`). Framing signals are active high (`LRFS = 0` and `LTFS = 0`).

Figure 12-31 through Figure 12-36 show framing for receiving data.

In Figure 12-31 and Figure 12-32, the normal framing mode is shown for non-continuous data (any number of TSCLK or RSCLK cycles between words) and continuous data (no TSCLK or SCLK cycles between words).

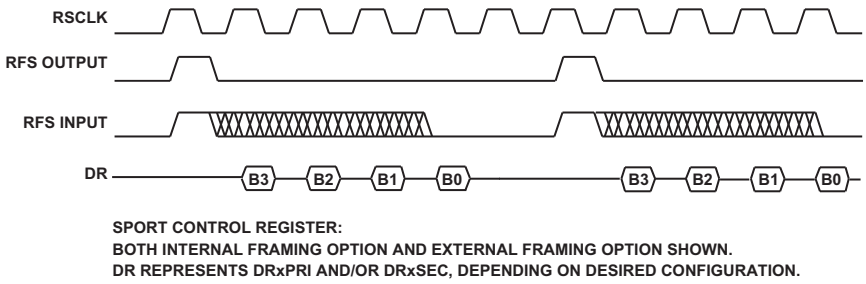

Figure 12-31. SPORT Receive, Normal Framing

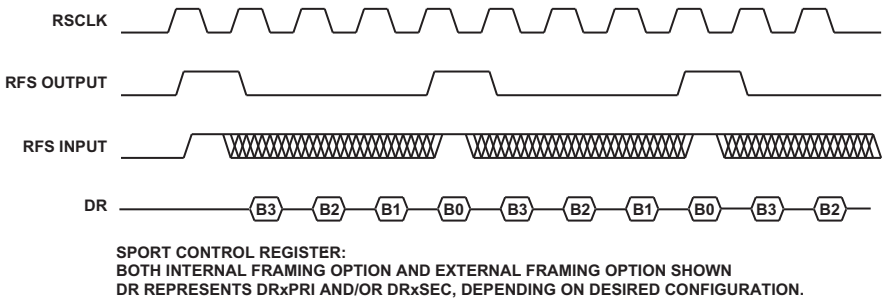

Figure 12-32. SPORT Continuous Receive, Normal Framing

Figure 12-33 and Figure 12-34 show non-continuous and continuous receiving in the alternate framing mode. These four figures show the input timing requirement for an externally generated frame sync and also the output timing characteristic of an internally generated frame sync. Note the output meets the input timing requirement; therefore, with two

SPORT channels used, one SPORT channel could provide RFS for the other SPORT channel.

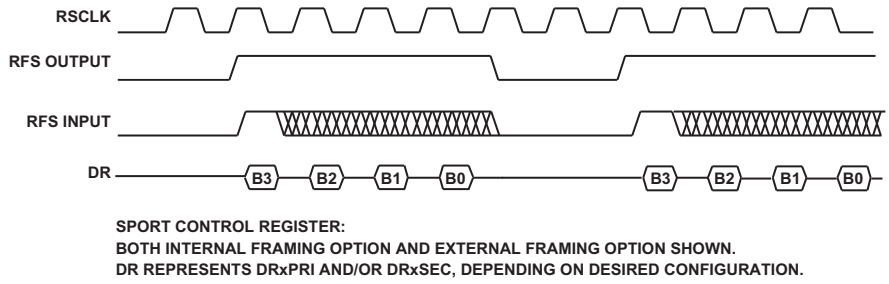

Figure 12-33. SPORT Receive, Alternate Framing

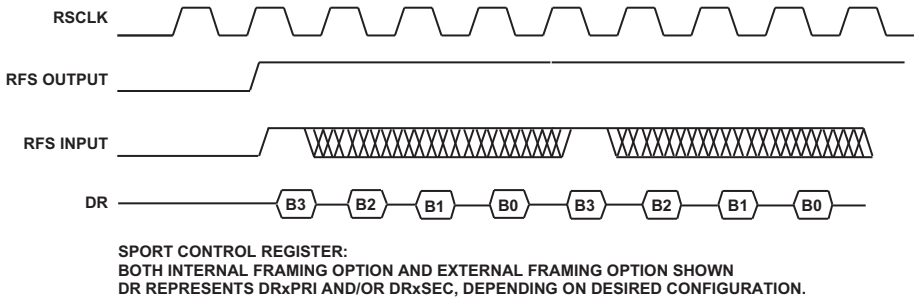

Figure 12-34. SPORT Continuous Receive, Alternate Framing

Figure 12-35 and Figure 12-36 show the receive operation with normal framing and alternate framing, in the unframed mode. A single frame sync signal occurs only at the start of the first word, either one RSCLK before the first bit (in normal mode) or at the same time as the first bit (in alternate

mode). This mode is appropriate for multi-word bursts (continuous reception).

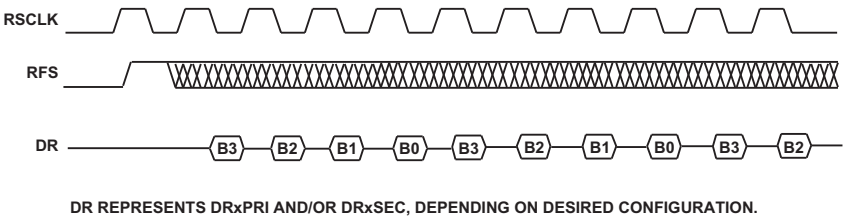

Figure 12-35. SPORT Receive, Unframed Mode, Normal Framing

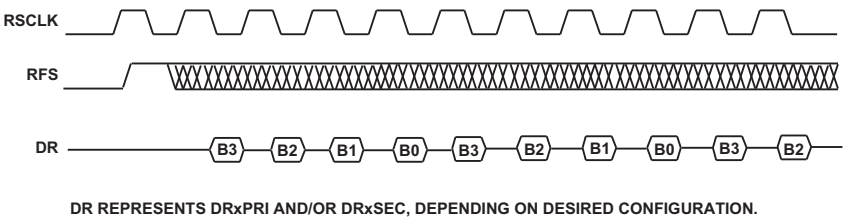

Figure 12-36. SPORT Receive, Unframed Mode, Alternate Framing

Figure 12-37 through Figure 12-42 show framing for transmitting data and are very similar to Figure 12-31 through Figure 12-36.

In Figure 12-37 and Figure 12-38, the normal framing mode is shown for non-continuous data (any number of TSCLK cycles between words) and continuous data (no TSCLK cycles between words). Figure 12-39 and Figure 12-40 show non-continuous and continuous transmission in the alternate framing mode. As noted previously for the receive timing diagrams, the RFS output meets the RFS input timing requirement.

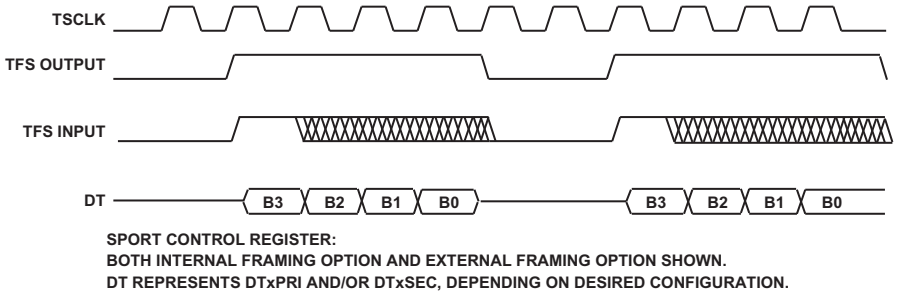

Figure 12-37. SPORT Transmit, Normal Framing

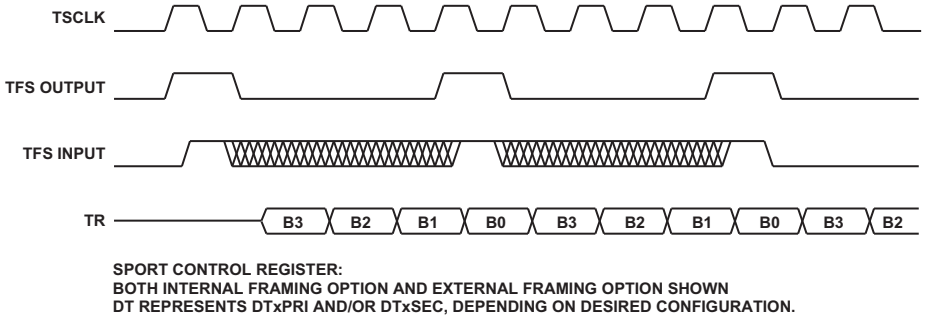

Figure 12-38. SPORT Continuous Transmit, Normal Framing

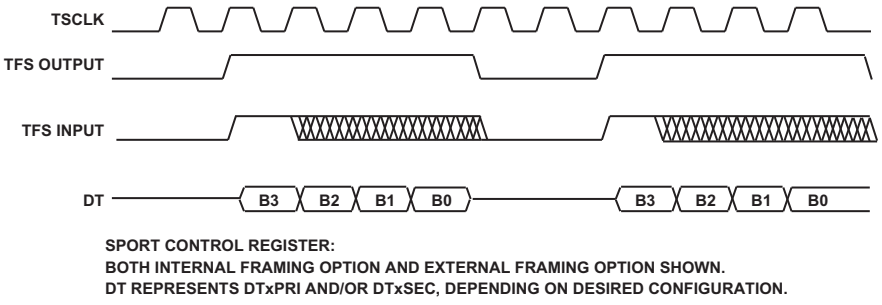

Figure 12-39. SPORT Transmit, Alternate Framing

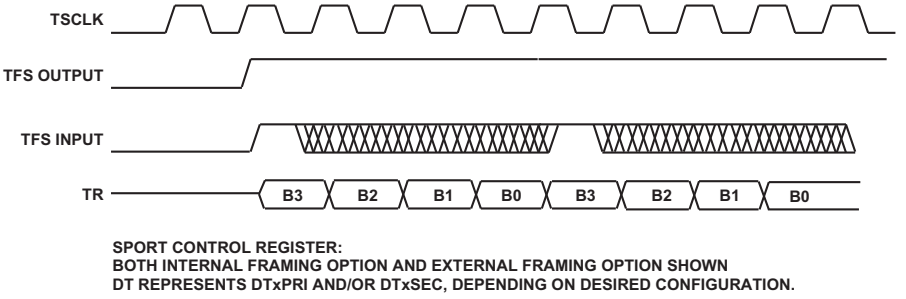

Figure 12-40. SPORT Continuous Transmit, Alternate Framing

Figure 12-41 and Figure 12-42 show the transmit operation with normal framing and alternate framing, respectively, in the unframed mode. A single frame sync signal occurs only at the start of the first word, either one TSCLK before the first bit (in normal mode) or at the same time as the first bit (in alternate mode).

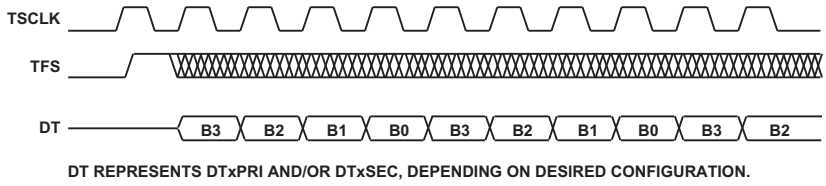

Figure 12-41. SPORT Transmit, Unframed Mode, Normal Framing

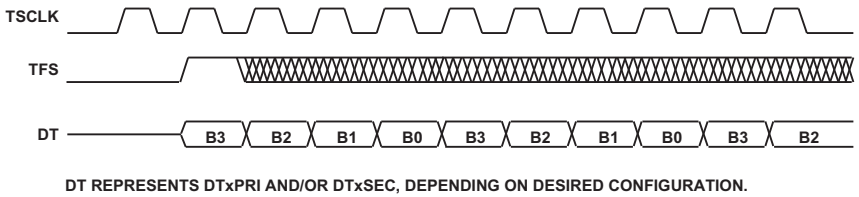

Figure 12-42. SPORT Transmit, Unframed Mode, Alternate Framing

13 UART PORT CONTROLLER

The Universal Asynchronous Receiver/Transmitter (UART) is a full-duplex peripheral compatible with PC-style industry standard UARTs. The UART converts data between serial and parallel formats. The serial communication follows an asynchronous protocol that supports various word length, stop bits, and parity generation options. The UART includes interrupt handling hardware. Interrupts can be generated from 12 different events.

The UART supports the half-duplex IrDA® (Infrared Data Association) SIR (9.6/115.2 Kbps rate) protocol. This is a mode-enabled feature.

Modem status and control functionality is not supported by the UART module, but may be implemented using General Purpose IO (GPIO) pins.

The UART is a DMA capable peripheral with support for separate TX and RX DMA master channels. It can be used in either DMA or programmed non-DMA mode of operation. The non-DMA mode requires software management of the data flow using either interrupts or polling. The DMA method requires minimal software intervention as the DMA engine itself moves the data. See Chapter 9, “Direct Memory Access” for more information on DMA.

Either one of the peripheral timers can be used to provide a hardware assisted autobaud detection mechanism for use with the UART. See Chapter 15, “Timers” for more information.

Serial Communications

The UART follows an asynchronous serial communication protocol with these options:

- 5 – 8 data bits
- 1, 1½, or 2 stop bits
- None, even, or odd parity
- Baud rate = `SCLK/(16 × Divisor)`, where `SCLK` is the system clock frequency and Divisor can be a value from 1 to 65536

All data words require a start bit and at least one stop bit. With the optional parity bit, this creates a 7- to 12-bit range for each word. The format of received and transmitted character frames is controlled by the Line Control register (`UART_LCR`). Data is always transmitted and received least significant bit (LSB) first.

Figure 13-1 shows a typical physical bit stream measured on the TX pin.

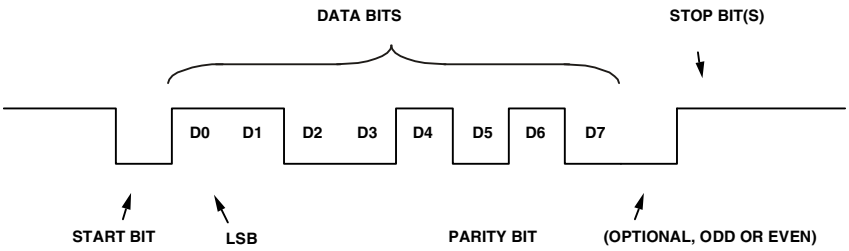

Figure 13-1. Bit Stream on the TX Pin

UART Control and Status Registers

The processor provides a set of PC-style industry standard control and status registers for each UART. These Memory Mapped Registers (MMR) are byte-wide registers that are mapped as half words with the most significant byte zero filled.

Consistent with industry standard interfaces, multiple registers are mapped to the same address location. The Divisor Latch registers (UART_DLH and UART_DLL) share their addresses with the Transmit Holding register (UART_THR), the Receive Buffer register (UART_RBR), and the Interrupt Enable register (UART_IER). The Divisor Latch Access bit (DLAB) in the Line Control Register (UART_LCR) controls which set of registers is accessible at a given time. Software must use 16-bit word load/store instructions to access these registers.

Transmit and receive channels are both buffered. The UART_THR register buffers the Transmit Shift register (TSR) and the UART_RBR register buffers the Receive Shift register (LSR). The Shift registers are not directly accessible by software.

UART Line Control Register (UART_LCR)

The UART Line Control register (UART_LCR) controls the format of received and transmitted character frames (Figure 13-2).

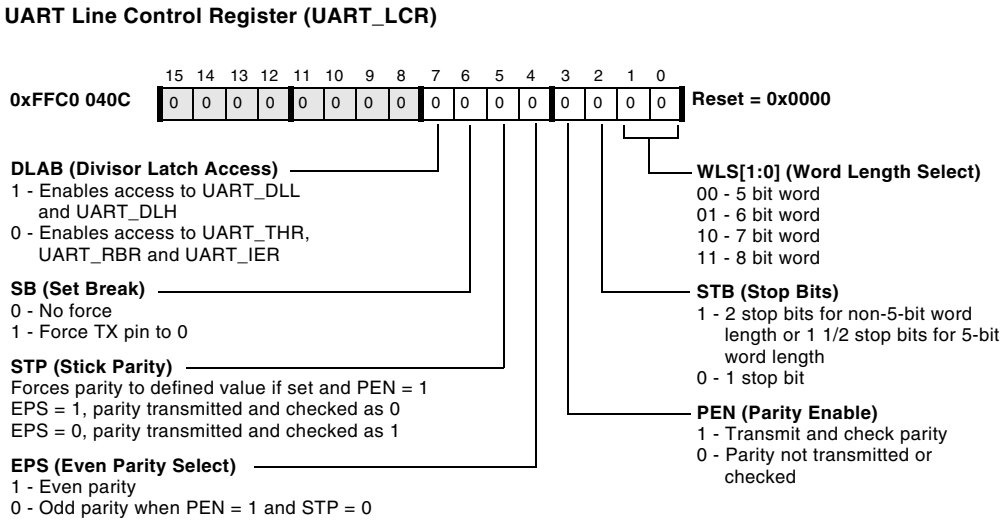

Figure 13-2. UART Line Control Register

The SB bit functions even when the UART clock is disabled. Since the TX pin normally drives high, it can be used as a flag output pin, if the UART is not used.

Modem Control Register (UART_MCR)

The Modem Control register `UART_MCR` controls the UART port, as shown in Figure 13-3. Even if modem functionality is not supported, the Modem Control register is available in order to support the loopback mode.

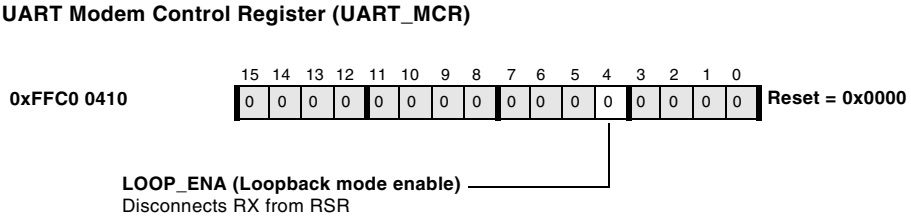

Figure 13-3. UART Modem Control Register

Loopback mode disconnects the receiver's input from the RX pin, but redirects it to the transmit output internally.

UART Line Status Register (UART_LSR)

The UART Line Status register contains UART status information as shown in Figure 13-4.

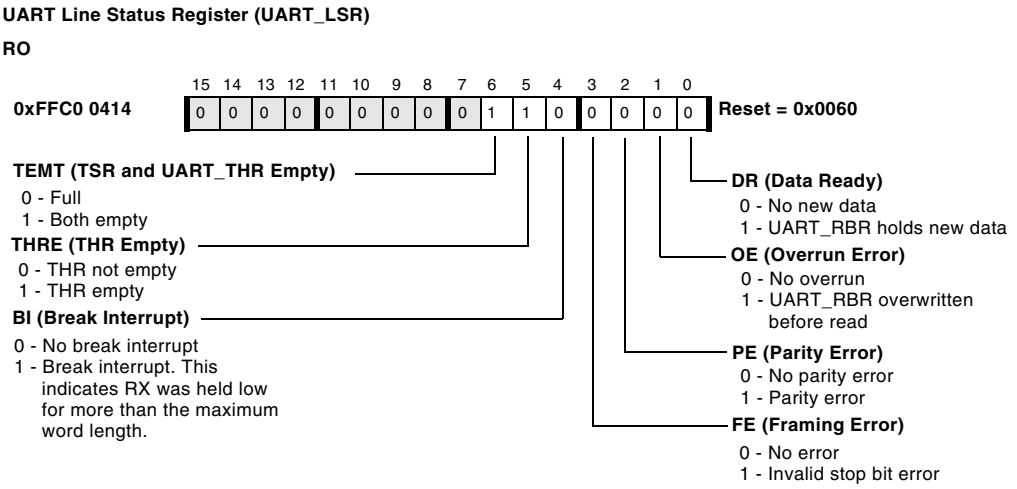

Figure 13-4. UART Line Status Register

The Break Interrupt (BI), Overrun Error (OE), Parity Error (PE) and Framing Error (FE) bits are cleared when the UART Line Status register (UART_LSR) is read. The Data Ready (DR) bit is cleared when the UART Receive Buffer register (UART_RBR) is read.

Because of the destructive nature of these read operations, special care should be taken. For more information, see "Speculative Load Execution" on page 6-70 and "Conditional Load Behavior" on page 6-71.

The THRE bit indicates that the UART transmit channel is ready for new data and software can write to UART_THR. Writes to UART_THR clear the THRE bit. It is set again when data is copied from the UART_THR to the Transmit

Shift register (TSR). The TEMT bit can be evaluated to determine whether a recently initiated transmit operation has been completed.

UART Transmit Holding Register (UART_THR)

A write to the UART Transmit Holding register (UART_THR) initiates the transmit operation (Figure 13-5). The data is moved to the internal Transmit Shift register (TSR) where it is shifted out at a baud rate equal to SCLK/(16 × Divisor) with start, stop, and parity bits appended as required. All data words begin with a 1-to-0-transition start bit. The transfer of data from UART_THR to the Transmit Shift register sets the Transmit Holding Register Empty (THRE) status flag in the UART Line Status register (UART_LSR).

The write-only UART_THR register is mapped to the same address as the read-only UART_RBR and UART_DLL. To access UART_THR, the DLAB bit in UART_LCR must be cleared. When the DLAB bit is cleared, writes to this address target the UART_THR register, and reads from this address return the UART_RBR register.

Note data is transmitted and received least significant bit first (LSB) followed by the most significant bits (MSBs).

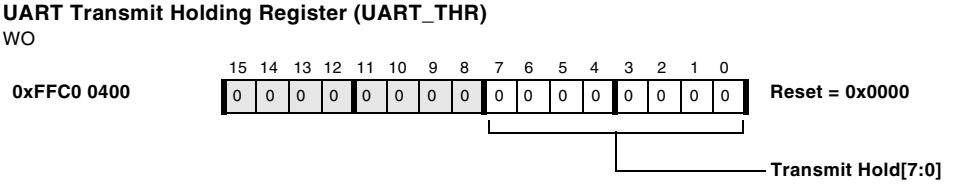

Figure 13-5. UART Transmit Holding Register

UART Receive Buffer Register (UART_RBR)

The receive operation uses the same data format as the transmit configuration, except that the number of stop bits is always assumed to be 1. After detection of the start bit, the received word is shifted into the Receive Shift register (RSR) at a baud rate of SCLK/(16 × Divisor). After the appropriate number of bits (including stop bit) is received, the data and any status are updated and the Receive Shift register is transferred to the UART Receive Buffer register (UART_RBR), shown in Figure 13-6. After the transfer of the received word to the UART_RBR buffer and the appropriate synchronization delay, the Data Ready (DR) status flag is updated.

A sampling clock equal to 16 times the baud rate samples the data as close to the midpoint of the bit as possible. Because the internal sample clock may not exactly match the asynchronous receive data rate, the sampling point drifts from the center of each bit. The sampling point is resynchronized with each start bit, so the error accumulates only over the length of a single word. A receive filter removes spurious pulses of less than two times the sampling clock period.

The read-only UART_RBR register is mapped to the same address as the write-only UART_THR and UART_DLL. To access UART_RBR, the DLAB bit in UART_LCR must be cleared. When the DLAB bit is cleared, writes to this address target the UART_THR register, while reads from this address return the UART_RBR register.

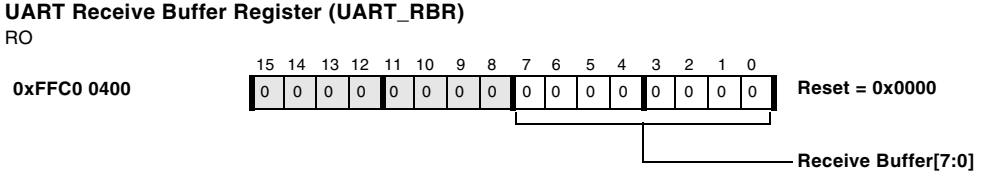

Figure 13-6. UART Receive Buffer Register

UART Interrupt Enable Register (UART_IER)

The UART Interrupt Enable register UART_IER is used to enable requests for system handling of empty or full states of UART data registers (Figure 13-7). Unless polling is used as a means of action, the ERBFI and/or ETBEI bits in this register are normally set.

Setting this register without enabling system DMA causes the UART to notify the processor of data inventory state by means of interrupts. For proper operation in this mode, system interrupts must be enabled, and appropriate interrupt handling routines must be present. For backward compatibility, the UART_IIR still reflects the correct interrupt status.

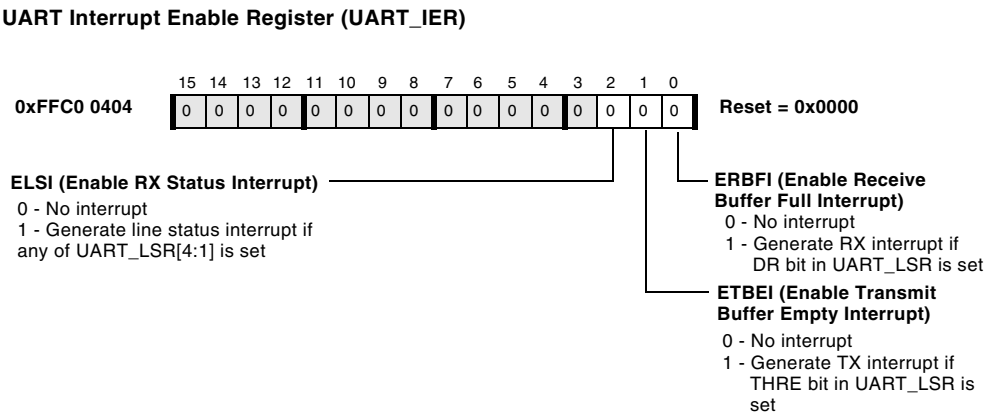

Figure 13-7. UART Interrupt Enable Register

The UART features three separate interrupt channels to handle data transmit, data receive and line status events independently, regardless whether DMA is enabled or not.

With system DMA enabled, the UART uses DMA to transfer data to or from the processor. Dedicated DMA channels are available to receive and transmit operation. Line error handling can be configured completely independently from the receive/transmit setup.

The UART_IER register is mapped to the same address as UART_DLH. To access UART_IER, the DLAB bit in UART_LCR must be cleared.

UART's DMA is enabled by first setting up the system DMA control registers and then enabling the UART ERBFI and ETBEI interrupts in the UART_IER register. This is because the interrupt request lines double as DMA request lines. Depending on whether DMA is enabled or not, upon receiving these requests, the DMA control unit either generates a direct memory access or passes the UART interrupt on to the system interrupt handling unit. However, UART's error interrupt goes directly to the system interrupt handling unit, bypassing the DMA unit completely.

The ELSI bit enables interrupt generation on an independent interrupt channel when any of the following conditions are raised by the respective bit in the UART Line Status register (UART_LSR):

- Receive Overrun Error (OE)
- Receive Parity Error (PE)
- Receive Framing Error (FE)
- Break Interrupt (BI)

When the ETBEI bit is set in the UART_IER register, the UART module immediately issues an interrupt or DMA request. When initiating the transmission of a string, no special handling of the first character is required. Set the ETBEI bit and let the interrupt service routine load the first character from memory and write it to the UART_THR register in the normal manner. Accordingly, the ETBEI bit should be cleared if the string transmission has completed.

UART Interrupt Identification Register (UART_IIR)

For legacy reasons, the UART Interrupt Identification register UART_IIR (Figure 13-8) still reflects the UART interrupt status. Legacy operation may require you to bundle all UART interrupt sources to a single

interrupt channel, servicing them all by the same software routine. This can be established by globally assigning all UART interrupts to the same interrupt priority, by using the system interrupt (SIC) controller.

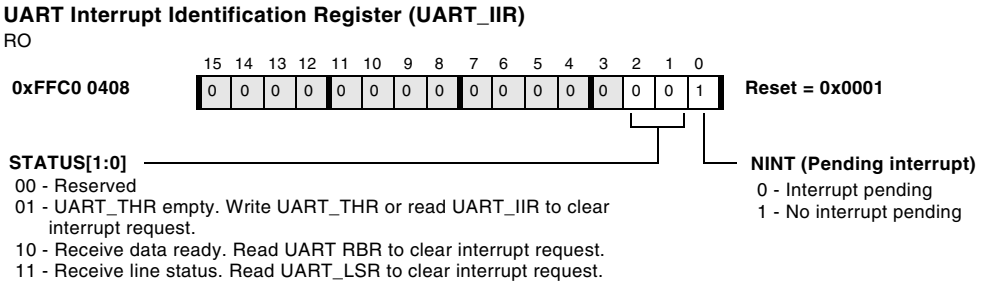

Figure 13-8. UART Interrupt Identification Register

When cleared, the Pending Interrupt bit (NINT) signals that an interrupt is pending. The STATUS field indicates the highest priority pending interrupt. The receive line status has the highest priority; the UART_THR empty interrupt has the lowest priority. In the case where both interrupts are signalling, the UART_IIR reads 0x06.

When a UART interrupt is pending, the interrupt service routine (ISR) needs to clear the interrupt latch explicitly. The following figure shows how to clear any of the three latches.

The TX interrupt request is cleared by writing new data to the UART_THR register or by reading the UART_IIR register. Please note the special role of the UART_IIR register read in the case where the service routine does not want to transmit further data.

If software stops transmission, it must read the UART_IIR register to reset the interrupt request. As long as the UART_IIR register reads 0x04 or 0x06 (indicating that another interrupt of higher priority is pending), the UART_THR empty latch cannot be cleared by reading UART_IIR.

If either the Line Status interrupt or the Receive Data interrupt has been assigned a lower interrupt priority by the SIC, a deadlock condition can occur. To avoid this, always assign the lowest priority of the enabled UART interrupts to the UART_THR empty event.

Because of the destructive nature of these read operations, special care should be taken. For more information, see "Speculative Load Execution" on page 6-70 and "Conditional Load Behavior" on page 6-71.

UART Divisor Latch Registers (UART_DLL, UART_DLH)

The bit rate is characterized by the system clock (SCLK) and the 16-bit Divisor. The Divisor is split into the Divisor Latch Low Byte register (UART_DLL) and the Divisor Latch High Byte register (UART_DLH). See Figure 13-9.

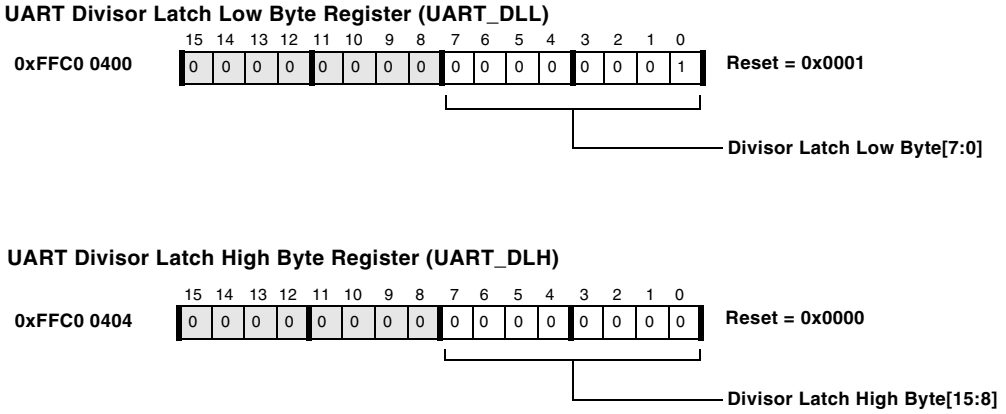

Figure 13-9. UART Divisor Latch Registers

These registers form a 16-bit divisor. The baud clock is divided by 16.

BAUD RATE = SCLK/(16 × Divisor)

Divisor = 65,536 when UART_DLL = UART_DLH = 0

The UART_DLL register is mapped to the same address as the UART_THR and UART_RBR registers. The UART_DLH register is mapped to the same address as the Interrupt Enable register (UART_IER). The DLAB bit in UART_LCR must be set before the UART Divisor Latch registers can be accessed.

Note the 16-bit divisor formed by UART_DLH and UART_DLL resets to 0x0001, resulting in the highest possible clock frequency by default. If the UART is not used, disabling the UART clock will save power. The UART_DLH and UART_DLL registers can be programmed by software before or after setting the UCEN bit.

Table 13-1 provides example divide factors required to support most standard baud rates.

Table 13-1. UART Baud Rate Examples with 100 MHz SCLK

Baud Rate	DL	Actual	% Error
2400	2604	2400.15	.006
4800	1302	4800.31	.007
9600	651	9600.61	.006
19200	326	19171.78	.147
38400	163	38343.56	.147
57600	109	57339.45	.452
115200	54	115740.74	.469
921600	7	892857.14	3.119
6250000	1	6250000	-

Careful selection of SCLK frequencies, that is, even multiples of desired baud rates, can result in lower error percentages.

UART Scratch Register (UART_SCR)

The contents of the 8-bit UART Scratch register (Figure 13-10) is reset to 0x00. It is used for general-purpose data storage and does not control the UART hardware in any way.

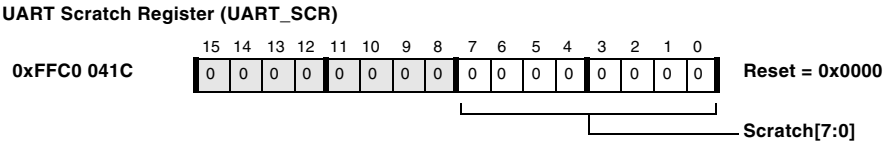

Figure 13-10. UART Scratch Register

UART Global Control Register (UART_GCTL)

The UART Global Control register UART_GCTL (Figure 13-11) contains the enable bit for internal UART clocks and for the IrDA mode of operation of the UART.

Please note that the UCEN bit was not present in previous UART implementations. It has been introduced to save power if the UART is not used. When porting code, be sure to enable this bit.

The IrDA TX Polarity Change bit and the IrDA RX Polarity Change bit are effective only in IrDA mode. The two force error bits, FPE and FFE, are intended for test purposes. They are useful for debugging software, especially in loopback mode.

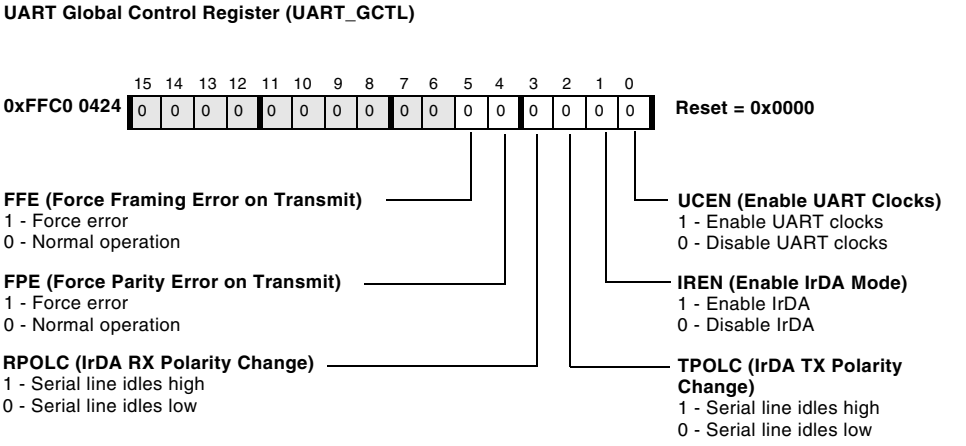

Figure 13-11. UART Global Control Register

Non-DMA Mode

In non-DMA mode, data is moved to and from the UART by the processor core. To transmit a character, load it into UART_THR. Received Data can be read from UART_RBR. The processor must write and read one character at a time.

To prevent any loss of data and misalignments of the serial datastream, the UART Line Status Register (UART_LSR) provides two status flags for handshaking: THRE and DR.

The THRE flag is set when UART_THR is ready for new data and cleared when the processor loads new data into UART_THR. Writing UART_THR when it is not empty overwrites the register with the new value and the previous character is never transmitted.

The DR flag signals when new data is available in UART_RBR. This flag is cleared automatically when the processor reads from UART_RBR. Reading UART_RBR when it is not full returns the previously received value. When

UART_RBR is not read in time, newly received data overwrites UART_RBR and the overrun (OE) flag is set.

With interrupts disabled, these status flags can be polled to determine when data is ready to move. Note that because polling is processor intensive, it is not typically used in real-time signal processing environments. Software can write up to two words into the UART_THR register before enabling the UART clock. As soon as the UCEN bit is set, those two words are sent.

Alternatively, UART writes and reads can be accomplished by Interrupt Service Routines (ISRs). Separate interrupt lines are provided for UART TX, UART RX, and UART Error. The independent interrupts can be enabled individually by the UART_IER register.

The ISRs can evaluate the Status bit field within the UART Interrupt Identification register (UART_IIR) to determine the signalling interrupt source. If more than one source is signalling, the status field displays the one with the highest priority. Interrupts also must be assigned and unmasked by the processor interrupt controller. The interrupt service routines must clear the interrupt latches explicitly. See Figure 13-8 on page 13-11.

DMA Mode

In this mode, separate receive (RX) and transmit (TX) DMA channels move data between the UART and memory. The software does not have to move data, it just has to set up the appropriate transfers either through the descriptor mechanism or through Autobuffer mode.

No additional buffering is provided in the UART DMA channel, so the latency requirements are the same as in non-DMA mode. However, the latency is determined by the bus activity and arbitration mechanism and not by the processor loading and interrupt priorities. For more information, see Chapter 9, “Direct Memory Access”.

DMA interrupt routines must explicitly write 1s to the corresponding DMA IRQ status registers explicitly to clear the latched request of the pending interrupt.

The UART's DMA is enabled by first setting up the system DMA control registers and then enabling the UART ERBFI and ETBEI interrupts in the UART_IER register. This is because the interrupt request lines double as DMA request lines. Depending on whether DMA is enabled or not, upon receiving these requests, the DMA control unit either generates a direct memory access or passes the UART interrupt on to the system interrupt handling unit. However, the UART's error interrupt goes directly to the system interrupt handling unit, bypassing the DMA unit completely.

The UART's DMA supports 8-bit operation.

Mixing Modes

Non-DMA and DMA modes use different synchronization mechanisms. Consequently, any serial communication must be complete before switching from non-DMA to DMA mode or vice-versa. In other words, before switching from non-DMA transmission to DMA transmission, make sure that both UART_THR and the internal Transmit Shift register (TSR) are empty by testing the THRE and the TEMT status bits in UART_LSR. Otherwise, the processor must wait until the 2-bit DMA Buffer Status field within the appropriate UART Transmit DMA Configuration register (UART_CONFIG_TX) is clear.

When switching from DMA to non-DMA operation, make sure that both the receive (RX) and transmit (TX) DMA channels have completely transferred their data, including data contained in the DMA FIFOs. While the DMA RX interrupt indicates the last data word has been written to memory (and has left the DMA FIFO), the DMA TX interrupt indicates the last data word has left memory (and has entered the DMA FIFO). The processor must wait until the TX FIFO is empty, by testing that the

DMA_RUN status bit in the TX channel's IRQ_STATUS register is clear, before it is safe to disable the DMA channel.

IrDA Support

Aside from the standard UART functionality, the UART also supports half-duplex serial data communication via infrared signals, according to the recommendations of the Infrared Data Association (IrDA). The physical layer known as IrDA SIR (9.6/115.2 Kbps rate) is based on return-to-zero-inverted (RZI) modulation. Pulse position modulation is not supported.

Using the 16x data rate clock, RZI modulation is achieved by inverting and modulating the non-return-to-zero (NRZ) code normally transmitted by the UART. On the receive side, the 16x clock is used to determine an IrDA pulse-sample window, from which the RZI-modulated NRZ code is recovered.

IrDA support is enabled by setting the IREN bit in the UART Global Control register. The IrDA application requires external transceivers.

IrDA Transmitter Description

To generate the IrDA pulse transmitted by the UART, the normal NRZ output of the transmitter is first inverted, so a 0 is transmitted as a high pulse of 16 UART clock periods and a 1 is transmitted as a low pulse for 16 UART clock periods. The leading edge of the pulse is then delayed by six UART clock periods. Similarly, the trailing edge of the pulse is truncated by eight UART clock periods. This results in the final representation of the original 0 as a high pulse of only 3/16 clock periods in a 16-cycle UART clock period. The pulse is centered around the middle of the bit time, as shown in Figure 13-12. The final IrDA pulse is fed to the off-chip infrared driver.

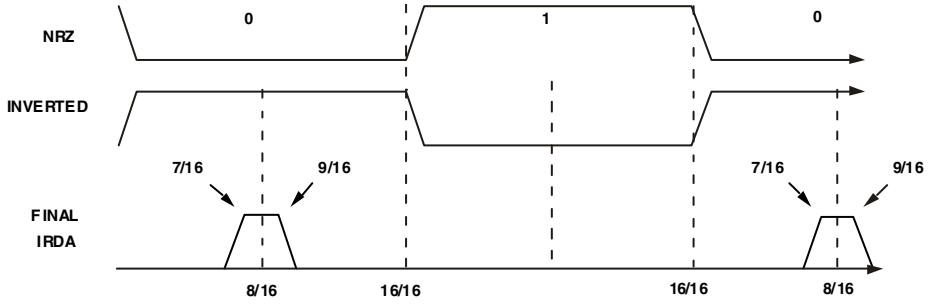

Figure 13-12. IrDA Transmit Pulse

This modulation approach ensures a pulse width output from the UART of three cycles high out of every 16 UART clock cycles. As shown in Table 13-1 on page 13-13, the error terms associated with the baud rate generator are very small and well within the tolerance of most infrared transceiver specifications.

IrDA Receiver Description

The IrDA receiver function is more complex than the transmit function. The receiver must discriminate the IrDA pulse and reject noise. To do this, the receiver looks for the IrDA pulse in a narrow window centered around the middle of the expected pulse.

Glitch filtering is accomplished by counting 16 system clocks from the time an initial pulse is seen. If the pulse is absent when the counter expires, it is considered a glitch. Otherwise, it is interpreted as a 0. This is acceptable because glitches originating from on-chip capacitive cross-coupling typically do not last for more than a fraction of the system clock period. Sources outside of the chip and not part of the transmitter can be avoided by appropriate shielding. The only other source of a glitch is the transmitter itself. The processor relies on the transmitter to perform within specification. If the transmitter violates the specification, unpredictable results may occur. The 4-bit counter adds an extra level of

protection at a minimal cost. Note because the system clock can change across systems, the longest glitch tolerated is inversely proportional to the system clock frequency.

The receive sampling window is determined by a counter, clocked at the 16x bit-time sample clock. The sampling window is resynchronized with each start bit by centering the sampling window around the start bit.

The polarity of receive data is selectable, using the IRPOL bit. Figure 13-13 gives examples of each polarity type.

- IRPOL = 0 assumes that the receive data input idles 0 and each active 1 transition corresponds to a UART NRZ value of 0.
- IRPOL = 1 assumes that the receive data input idles 1 and each active 0 transition corresponds to a UART NRZ value of 0.

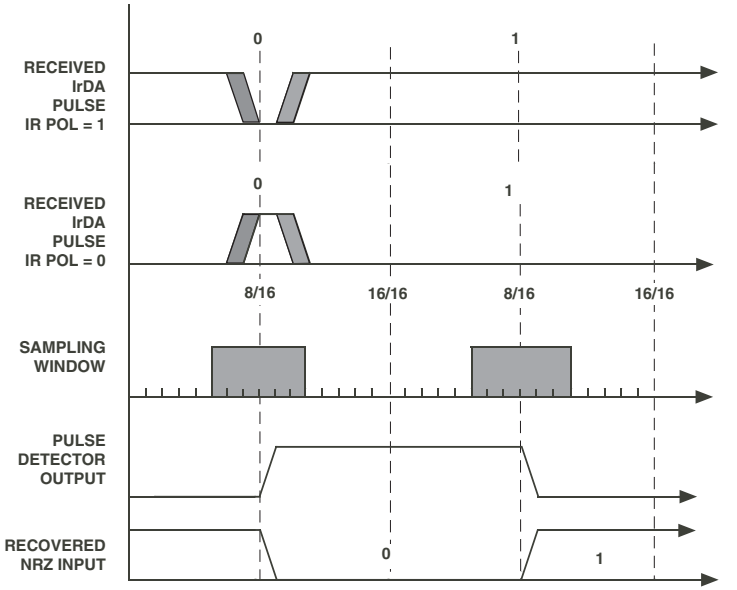

Figure 13-13. IrDA Receiver Pulse Detection

14 PROGRAMMABLE FLAGS

The processor supports 48 bidirectional programmable flags (PFx) or general purpose I/O pins, PF[0:47]. Each pin can be individually configured as either an input or an output by using the Flag Direction register (FIOn_DIR). When configured as an output, the Flag Data register (FIOn_FLAG_D) can be directly written to specify the state of all PFx pins. When configured as an output, the state written to the Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) registers determines the state driven by the output PFx pin. Regardless of how the pins are configured, as an input or an output, reading any of these FIOn_FLAG_* registers returns the state of each pin.

Each PFx pin can be configured to generate an interrupt. When a PFx pin is configured as an input, an interrupt can be generated according to the state of the pin (either high or low), an edge transition (low to high or high to low), or on both edge transitions (low to high *and* high to low). Input sensitivity is defined on a per-bit basis by the Flag Polarity register (FIOn_POLAR), the Flag Interrupt Sensitivity register (FIOn_EDGE) and the Flag Set on Both Edges register (FIOn_BOTH). Output sensitivity is defined on a per-bit basis by the Flag Polarity register. When a PFx pin is configured as an output (via FIOn_DIR) and the PFx inputs are enabled (via FIOn_INEN), enabling interrupts for the pin allows an interrupt to be generated by setting the PFx pin.

The processor provides two independent interrupt channels for the PFx pins. Identical in functionality, these are called Interrupt A and Interrupt B. Each interrupt channel has four mask registers associated with it, a Flag Interrupt Mask Data register (FIOn_MASKx_D), a Flag Interrupt Mask Set register (FIOn_MASKx_S) a Flag Interrupt Mask Clear register

(FIOn_MASKx_C), and a Flag Interrupt Mask Toggle register (FIOn_MASKx_T).

Each PFx pin is represented by a bit in each of these eight register sets. Writing a 1 to a bit in a Mask Set register enables interrupt generation for that PFx pin, while writing a 1 to a bit in a Mask Clear register disables interrupt generation for that PFx pin.

The interrupt masking can be toggled by writing a 1 to a bit in the Mask Toggle register. Additionally, the mask bits can be directly written by writing to the Mask Data register. This flexible mechanism allows each bit to generate Flag Interrupt A, Flag Interrupt B, both Flag Interrupts A and B, or neither.

When a PF pin is not used in a system, the input buffer can be disabled so that no external pullups or pulldowns are required on the unused pins. By default, the input buffers are disabled. They can be enabled via bits in the Flag Input Enable register (FIOn_INEN).

The PFx pins are multiplexed for use by the Parallel Peripheral Interface (PPI), Timers, Serial Peripheral Interface (SPI), SPORT, and UART.

The PFx pins are multiplexed for use by the Parallel Peripheral Interface (PPI), Timers, and Serial Peripheral Interface (SPI). Table 14-2 on page 14-5 shows the programmable flag pins and their multiplexed functionality.

The PF[47:40] pins are “mux’ed” with the upper eight data bits of PPI0 and the PF[39:32] pins are “mux’ed” with the upper eight data bits of PPI1. These pins are only available as programmable flags when they are not being used by the relevant PPI.

The PF[31:16] pins are “mux’ed” with the UART and both Serial Ports. These pins are only available as programmable flags under the following conditions:

- PF[29:28, 20:16] are available when Sport 0 is not active
- PF[31:30, 25:21] are available when Sport 1 is not active
- PF[27:26] are available when the UART is not active

The PF[7:0] pins are “mux’ed” with the SPI and Timer7:0. Table 14-1 shows the programmable flag “mux’ed” and the conditions when they are active.

Table 14-1. Programmable Flag and Peripheral MUX Conditions

PF Pin	“mux’ed” With	PF Active When
PF[47]	PPI0_D[15]	PPI0 not active or data length < 16
PF[46]	PPI0_D[14]	PPI0 not active or data length < 15
PF[45]	PPI0_D[13]	PPI0 not active or data length < 14
PF[44]	PPI0_D[12]	PPI0 not active or data length < 13
PF[43]	PPI0_D[11]	PPI0 not active or data length < 12
PF[42]	PPI0_D[10]	PPI0 not active or data length < 11
PF[41]	PPI0_D[9]	PPI0 not active or data length < 10
PF[40]	PPI0_D[8]	PPI0 not active or data length < 9
PF[39]	PPI1_D[15]	PPI1 not active or data length < 16
PF[38]	PPI1_D[14]	PPI1 not active or data length < 15
PF[37]	PPI1_D[13]	PPI1 not active or data length < 14
PF[36]	PPI1_D[12]	PPI1 not active or data length < 13
PF[35]	PPI1_D[11]	PPI1 not active or data length < 12
PF[34]	PPI1_D[10]	PPI1 not active or data length < 11
PF[33]	PPI1_D[9]	PPI1 not active or data length < 10
PF[32]	PPI1_D[8]	PPI1 not active or data length < 9

Table 14-1. Programmable Flag and Peripheral MUX Conditions (Cont'd)

PF Pin	"mux'ed" With	PF Active When
PF[31]	TSCLK1	SPORT1 not active
PF[30]	RSCLK1	SPORT1 not active
PF[29]	TSCLK0	SPORT0 not active
PF[28]	RSCLK0	SPORT0 not active
PF[27]	RX	UART not active
PF[26]	TX	UART not active
PF[25]	DR1SEC	SPORT1 not active
PF[24]	RFS1	SPORT1 not active
PF[23]	DT1PRI	SPORT1 not active
PF[22]	DT1SEC	SPORT1 not active
PF[21]	TFS1	SPORT1 not active
PF[20]	DR0SEC	SPORT0 not active
PF[19]	RFS0	SPORT0 not active
PF[18]	DT0PRI	SPORT0 not active
PF[17]	DT0SEC	SPORT0 not active
PF[16]	TFS0	SPORT0 not active
PF[15]	EXT CLK	Input to Timer 0-11
PF[14]		
PF[13]		
PF[12]		
PF[11]		
PF[10]		
PF[9]		
PF[8]		
PF[7]	SPISEL7 or TMR7	Timer 7 active or input for SPISEL7
PF[6]	SPISEL6 or TMR6	Timer 6 active or input for SPISEL6

Table 14-1. Programmable Flag and Peripheral MUX Conditions (Cont'd)

PF Pin	"mux'ed" With	PF Active When
PF[5]	SPISEL5 or TMR5	Timer 5 active or input for SPISEL5
PF[4]	SPISEL4 or TMR4	Timer 4 active or input for SPISEL4
PF[3]	SPISEL3 or TMR3	Timer 3 active or input for SPISEL3
PF[2]	SPISEL2 or TMR2	Timer 2 active or input for SPISEL2
PF[1]	SPISEL1 or TMR1	Timer 1 active or input for SPISEL1
PF[0]	$\overline{\text{SPISS}}$ or TMR0	Timer 0 active or output for $\overline{\text{SPISS}}$

Table 14-2. Programmable Flag Pins and Functionality

PF Pin	Peripheral That Shares the PF Pin		
	PPI	SPI	Timers7:0
0		Slave Select Input ($\overline{\text{SPISS}}$)	
1		Slave Select Enable 1 (SPISEL1)	Input clock
2		Slave Select Enable 2 (SPISEL2)	
3	Frame Sync 3	Slave Select Enable 3 (SPISEL3)	
4	I/O #15	Slave Select Enable 4 (SPISEL4)	
5	I/O #14	Slave Select Enable 5 (SPISEL5)	
6	I/O #13	Slave Select Enable 6 (SPISEL6)	
7	I/O #12	Slave Select Enable 7 (SPISEL7)	
8	I/O #11		
9	I/O #10		
10	I/O #9		

Table 14-2. Programmable Flag Pins and Functionality (Cont'd)

PF Pin	Peripheral That Shares the PF Pin		
	PPI	SPI	Timers7:0
11	I/O #8		
12	I/O #7		
13	I/O #6		
14	I/O #5		
15	I/O #4		

Table 14-3 describes how to use the peripheral function that shares the PF pin.

Table 14-3. Using Peripheral Function That Shares the PF Pin

PF Pin	To Use the Peripheral Function That Shares the PF Pin… (Assuming Peripheral is Enabled)		
	PPI	SPI	Timers7:0
0		Write '1' to PSSE in SPI_CTL	
1		Write '1' to FLS1 in SPI_FLG	Write '1' to CLK_SEL in TIMERx_ CONFIG
2		Write '1' to FLS2 in SPI_FLG	
3	Write '01' to PORT_CFG in PPIx_CTL (if PORT_DIR = '1'), or write '10' to PORT_CFG (if PORT_DIR = '0')	Write '1' to FLS3 in SPI_FLG	
4	Write '111' to DLEN in PPIx_CTL	Write '1' to FLS4 in SPI_FLG	
5	Write '110' to DLEN in PPIx_CTL	Write '1' to FLS5 in SPI_FLG	

Table 14-3. Using Peripheral Function That Shares the PF Pin (Cont'd)

PF Pin	To Use the Peripheral Function That Shares the PF Pin… (Assuming Peripheral is Enabled)		
	PPI	SPI	Timers7:0
6	Write '101' to DLEN in PPIx_CTL	Write '1' to FLS6 in SPI_FLG	
7	Write '100' to DLEN in PPIx_CTL	Write '1' to FLS7 in SPI_FLG	
8	Write '011' to DLEN in PPIx_CTL		
9	Write '010' to DLEN in PPIx_CTL		
10	Write '001' to DLEN in PPIx_CTL		
11	Write '000' to DLEN in PPIx_CTL		
12	Always enabled when PPI enabled		
13	Always enabled when PPI enabled		
14	Always enabled when PPI enabled		
15	Always enabled when PPI enabled		

For more details, please refer to the applicable datasheet, and the following chapters in this manual: Chapter 10, "SPI Compatible Port Controllers", Chapter 11, "Parallel Peripheral Interface", and Chapter 15, "Timers".

Programmable Flag Memory-Mapped Registers (MMRs)

The programmable flag (PFx) registers are part of the system memory-mapped registers (MMRs). The addresses of the programmable flag MMRs appear in Appendix B. Core access to the Flag Configuration registers is through the system bus.

Flag Direction Register (FIOn_DIR)

The Flag Direction register is a read-write register. Each bit position corresponds to a PFx pin. A logic 1 configures a PFx pin as an output, driving the state contained in the FIOn_FLAG_D register. A logic 0 configures a PFx pin as an input. The reset value of this register is 0x0000, making all PF pins inputs upon reset.

Note that when using the PFx pin as an input, the corresponding bit should also be set in the Flag Input Enable register.

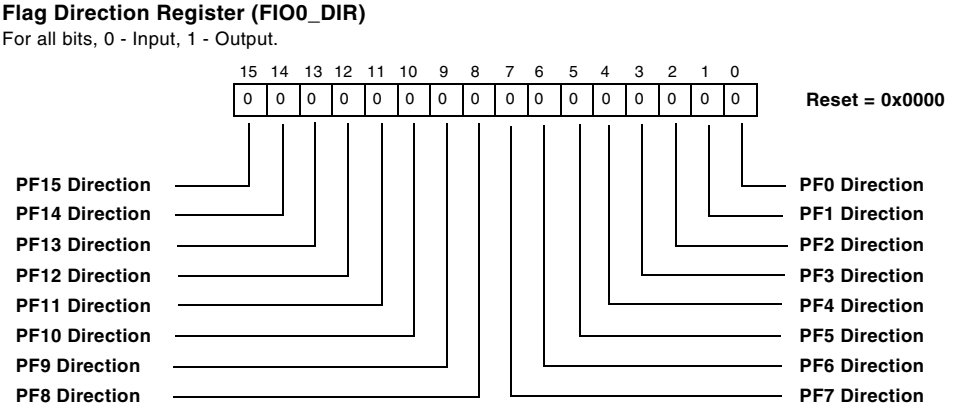

Figure 14-1. Flag (PF15 – PF0) Direction Register

Flag Direction Register (FIO1_DIR)
For all bits, 0 - Input, 1 - Output.

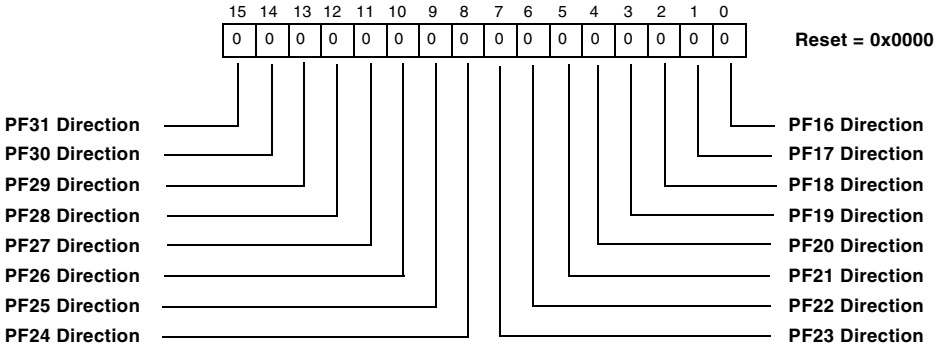

Figure 14-2. Flag (PF31 – PF16) Direction Register

Flag Direction Register (FIO2_DIR)
For all bits, 0 - Input, 1 - Output.

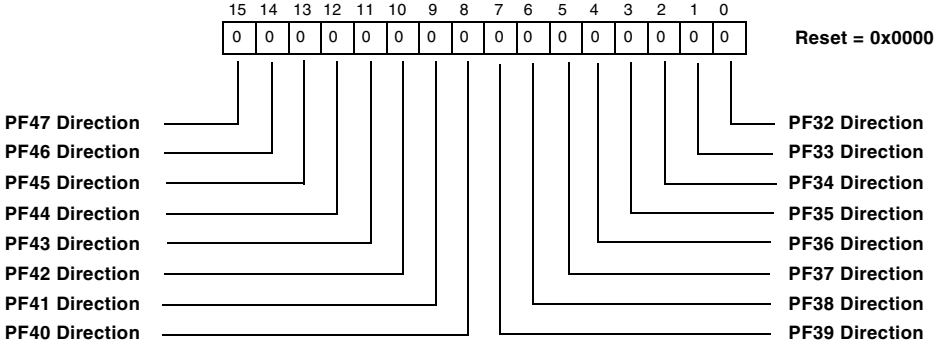

Figure 14-3. Flag (PF47 – PF32) Direction Register

Flag Value Registers

The processor has four Flag Value registers:

- Flag Data register (FIOn_FLAG_D)
- Flag Set register (FIOn_FLAG_S)
- Flag Clear register (FIOn_FLAG_C)
- Flag Toggle Direct register (FIOn_FLAG_T)

These registers are used to:

- Sense the value of the PFx pins defined as inputs
- Specify the state of PFx pins defined as outputs
- Clear interrupts generated by the PFx pins

Each PFx pin is represented by a bit in each of the four registers.

Reading any of the Flag Data, Flag Set, Flag Clear, or Flag Toggle registers returns the value of the PFx pins. The value returned shows the state of the PFx pins defined as outputs and the sense of PFx pins defined as inputs, based on the polarity and sensitivity settings of each pin.

Reading the Flag Data, Flag Set, Flag Clear, or Flag Toggle register after reset —results in 0x0000, because although the pins are inputs, the input buffers are not enabled. See Table 14-4 for guidance on how to interpret a value read from one of these registers, based on the settings of the FIOn_POLAR, FIOn_EDGE, and FIOn_BOTH registers.

For pins configured as edge-sensitive, a readback of 1 from one of these registers is sticky. That is, once it is set it remains set until cleared by user code. For level-sensitive pins, the pin state is checked every cycle, so the readback value will change when the original level on the pin changes.

Table 14-4. Flag Value Register Pin Interpretation

FIOn_POLAR	FIOn_EDGE	FIOn_BOTH	Effect of MMR Settings
0	0	X	Pin that is high reads as 1; pin that is low reads as 0
0	1	0	If rising edge occurred, pin reads as 1; otherwise, pin reads as 0
1	0	X	Pin that is low reads as 1; pin that is high reads as 0
1	1	0	If falling edge occurred, pin reads as 1; otherwise, pin reads as 0
X	1	1	If any edge occurred, pin reads as 1; otherwise, pin reads as 0

For more information about Flag Set, Flag Clear, and Flag Toggle registers, see "Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13.

Flag Data Register (FIOn_FLAG_D)

When written, the Flag Data register (Figure 14-4) directly specifies the state of all PFx pins. When read, the register returns the value of the PFx pins.

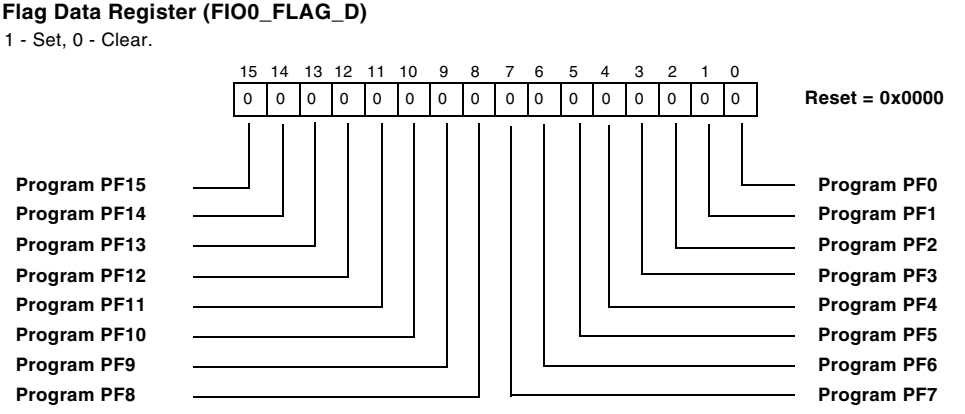

Figure 14-4. Flag (PF15 – PF0) Data Register

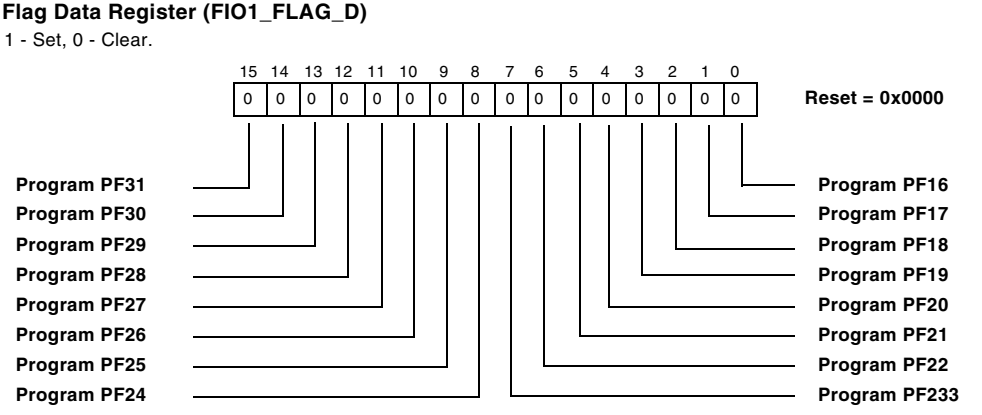

Figure 14-5. Flag (PF31 – PF16) Data Register

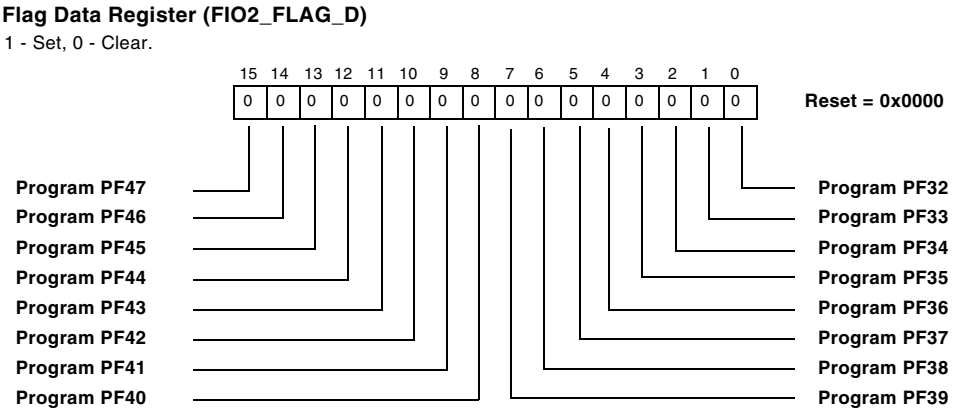

Figure 14-6. Flag (PF47 – PF32) Data Register

Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers

The Flag Set, Flag Clear, and Flag Toggle registers are used to:

- Set, clear or toggle the output state associated with each output PFx pin
- Clear the latched interrupt state captured from each input PFx pin

This mechanism is used to avoid the potential issues with more traditional read-modify-write mechanisms. Reading any of the these registers shows the flag pin state.

Figure 14-7 and Figure 14-10 represent the Flag Set and Flag Clear registers, respectively.

Flag Set Register (FIO0_FLAG_S)

Write-1-to-set.

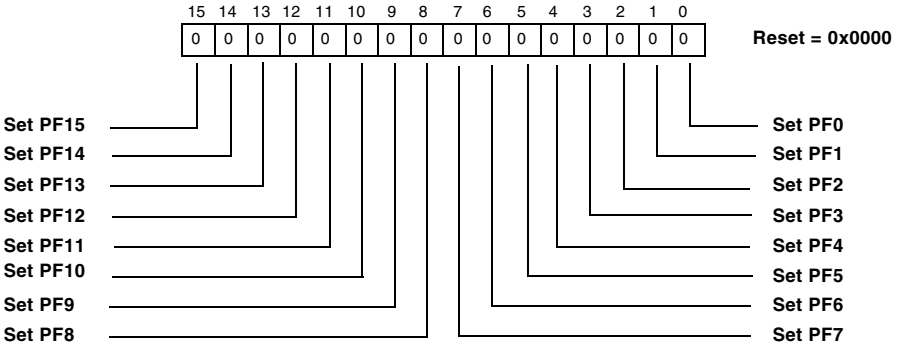

Figure 14-7. Flag (PF15 – PF0) Set Register

Flag Set Register (FIO1_FLAG_S)

Write-1-to-set.

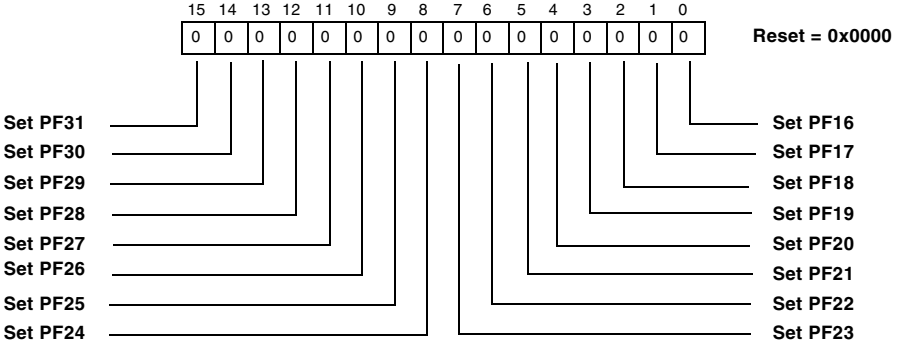

Figure 14-8. Flag (PF31 – PF16) Set Register

Flag Set Register (FIO2_FLAG_S)
Write-1-to-set.

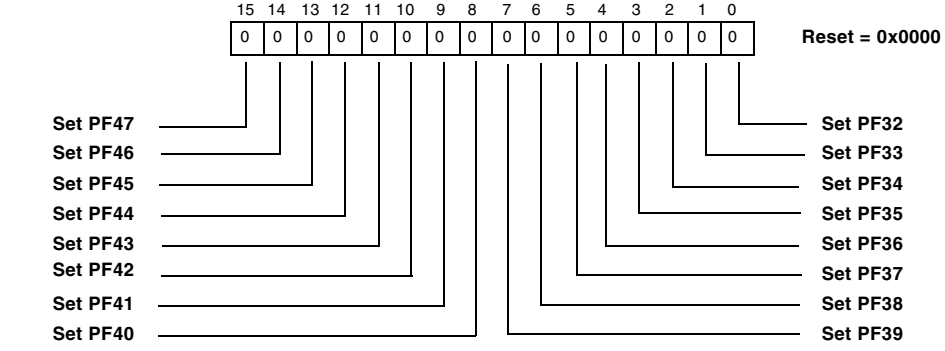

Figure 14-9. Flag (PF47 – PF32) Set Register

Flag Clear Register (FIO0_FLAG_C)
Write-1-to-clear.

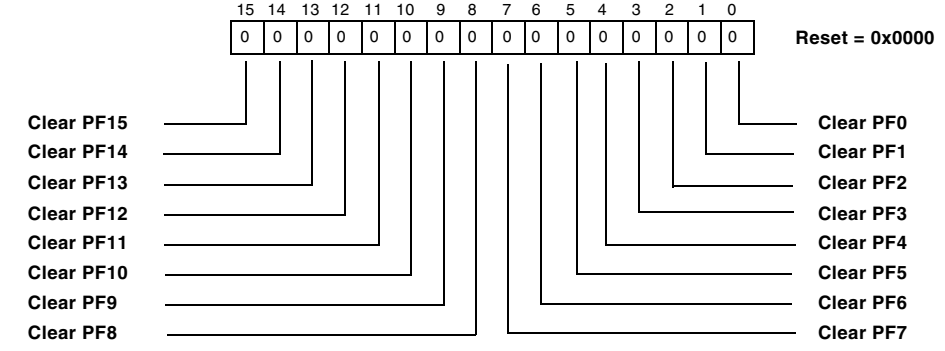

Figure 14-10. Flag (PF15 – PF0) Clear Register

Flag Clear Register (FIO1_FLAG_C)
Write-1-to-clear.

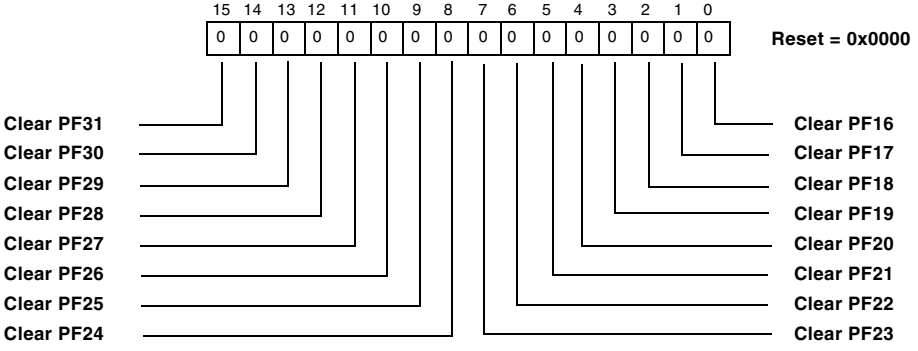

Figure 14-11. Flag (PF31 – PF16) Clear Register

Flag Clear Register (FIO2_FLAG_C)
Write-1-to-clear.

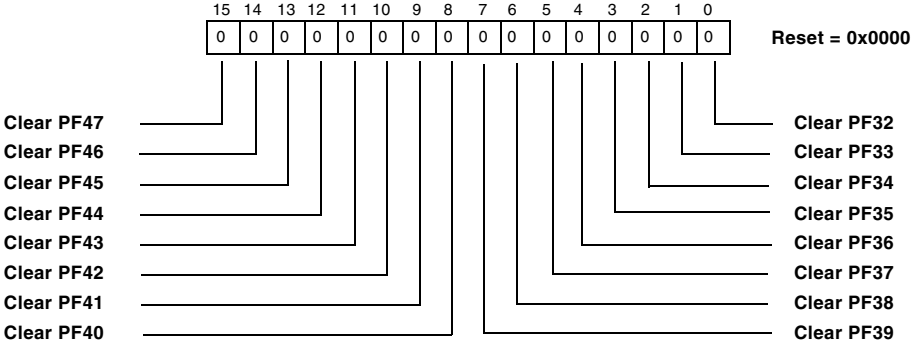

Figure 14-12. Flag (PF47 – PF32) Clear Register

Figure 14-13 represents the Flag Toggle register.

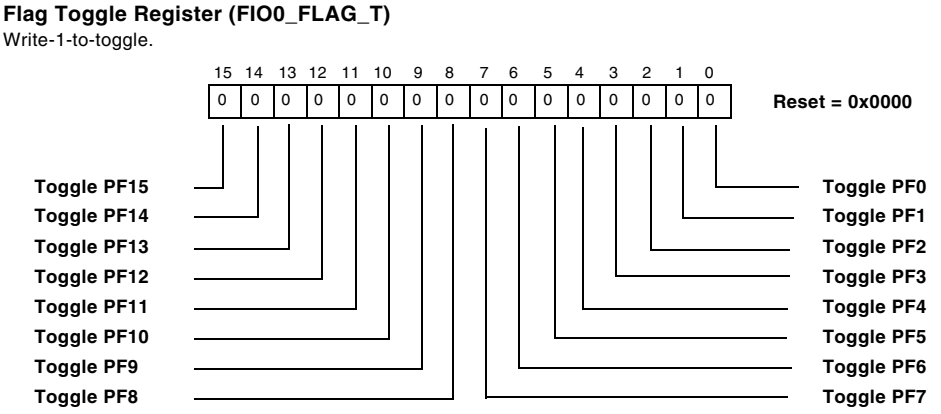

Figure 14-13. Flag (PF15 – PF0) Toggle Register

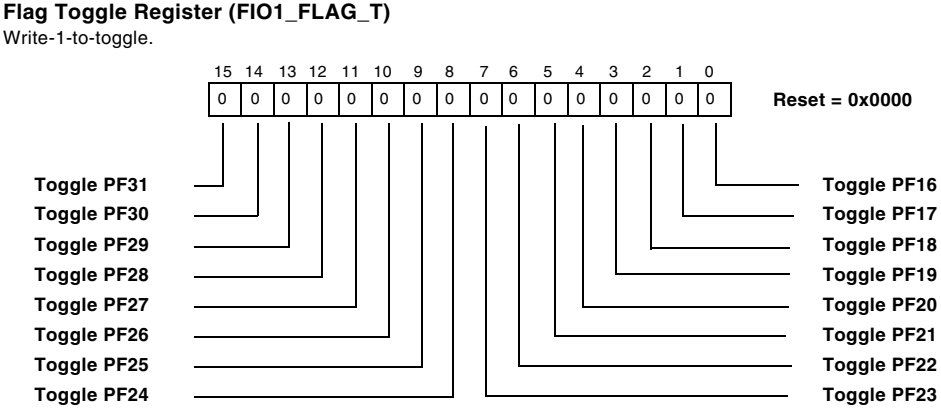

Figure 14-14. Flag (PF31 – PF16) Toggle Register

Flag Toggle Register (FIO2_FLAG_T)
Write-1-to-toggle.

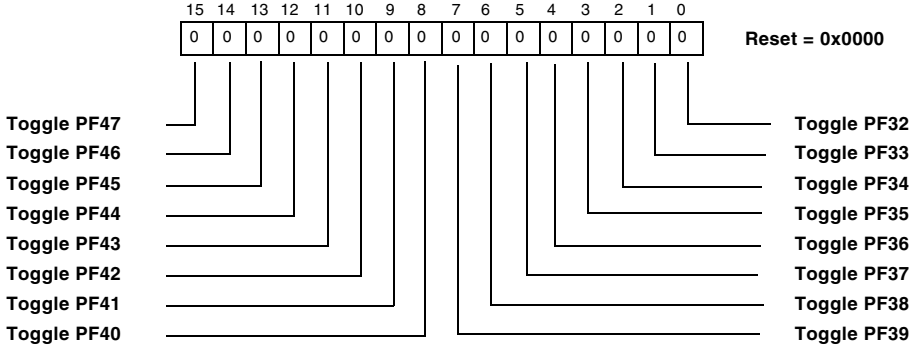

Figure 14-15. Flag (PF47 – PF32) Toggle Register

As an example of how these registers work, assume that PF[0] is configured as an output. Writing 0x0001 to the Flag Set register drives a logic 1 on the PF[0] pin without affecting the state of any other PFx pins. Writing 0x0001 to the Flag Clear register drives a logic 0 on the PF[0] pin without affecting the state of any other PFx pins. Writing a 0x0001 to the Flag Toggle register changes the pin state on PF[0] from logic zero to logic one or from logic one to logic zero, depending upon the existing pin state.

Writing a 0 to the Flag Set, Flag Clear, or Flag Toggle register has no effect on the value of the flag pin, and is ignored.

Reading the Flag Set or Flag Clear register shows:

- 0's for PFx pins defined as outputs and driven low
- 1's for pins (including PF[0] in the example above) defined as outputs and driven high, and
- The present sense of PFx pins defined as inputs

Input sense is based on FIOn_POLAR and FIOn_EDGE settings, as well as the logic level at each pin.

Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)

The Flag Mask Interrupt registers are implemented as complementary pairs of write-1-to-set, write-1-to-clear, and write-1-to-toggle registers. This implementation provides the ability to enable or disable a PFx pin to act as an interrupt without requiring read-modify-write accesses—or to directly specify the mask value with the data register.

Both Flag Interrupt A and Flag Interrupt B are supported by a set of four dedicated registers:

- Flag Mask Interrupt Data register
- Flag Mask Interrupt Set register
- Flag Mask Interrupt Clear register
- Flag Interrupt Toggle register

For diagrams of the registers that support Flag Interrupt A, see "Flag Mask Interrupt A Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S, FIOn_MASKA_T)" on page 14-22.

For diagrams of the registers that support Flag Interrupt B, see "Flag Mask Interrupt B Registers (FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-28.

Each PFx pin is represented by a bit in each of the eight registers. Table 14-5 shows the effect of writing 1 to a bit in a Mask Set, Mask Clear, or Mask Toggle register.

Table 14-5. Effect of Writing 1 to a Bit

Register	Effect of Writing 1 to a Bit in the Register
Mask Set	Enables interrupt generation for that PFx pin
Mask Clear	Disables interrupt generation for that PFx pin
Mask Toggle	Changes the state of interrupt generation capability
Mask Data	Enables interrupt generation for that PFx pin

Reading any of the [A and B] mask data, set, clear, or toggle registers returns the value of the current mask [A and B] data.

Interrupt A and Interrupt B operate independently. For example, writing 1 to a bit in the Flag Mask Interrupt A Set register does not affect Flag Interrupt B. This facility allows PFx pins to generate Flag Interrupt A, Flag Interrupt B, both Flag Interrupts A and B, or neither.

Note that a Flag Interrupt is generated by a logical OR of all unmasked PF pins for that interrupt. For example, if PF[0] and PF[1] are both unmasked for Flag Interrupt A, Flag Interrupt A will be generated when triggered by PF[0] or PF[1].

When using either rising or falling edge-triggered interrupts, the interrupt condition must be cleared each time a corresponding interrupt is serviced, by writing a “1” to the appropriate FIOn_FLAG_C bit.

At reset, all interrupts are masked.

Flag Interrupt Generation Flow

Figure 14-16 shows the process by which a Flag Interrupt A or a Flag Interrupt B event can be generated. Note that the flow is shown for only

one programmable flag, FlagN. However, a Flag Interrupt is generated by a logical OR of all unmasked PFx pins for that interrupt. For example, if only PF[0] and PF[1] are unmasked for Flag Interrupt A, this interrupt is generated when triggered by either PF[0] or PF[1].

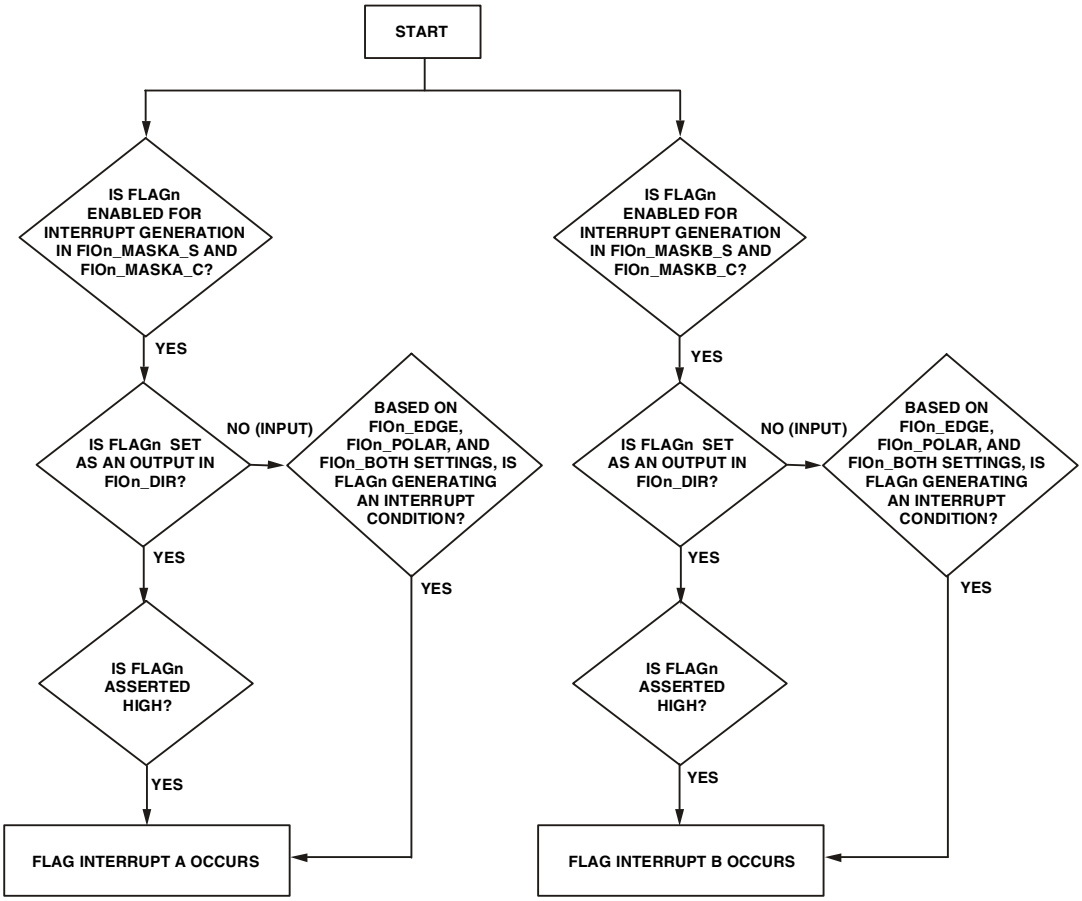

Figure 14-16. Flag Interrupt Generation Flow

Flag Mask Interrupt A Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S, FIOn_MASKA_T)

The register sets in Figure 14-17 to Figure 14-28 on page 14-27 support Flag Interrupt A. Details can be found on page 14-19.

Flag Mask Interrupt A Data Register (FIO0_MASKA_D)
For all bits, 1 - Enable.

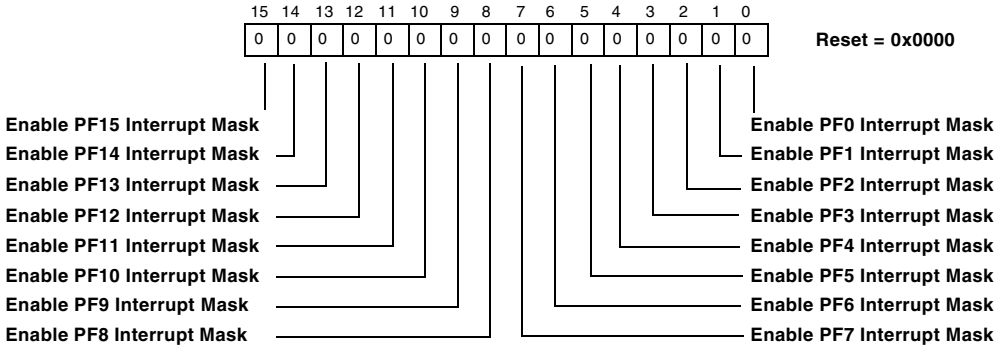

Figure 14-17. Flag (PF15 – PF0) Mask Interrupt A Data Register

Flag Mask Interrupt A Data Register (FIO1_MASKA_D)
For all bits, 1 - Enable.

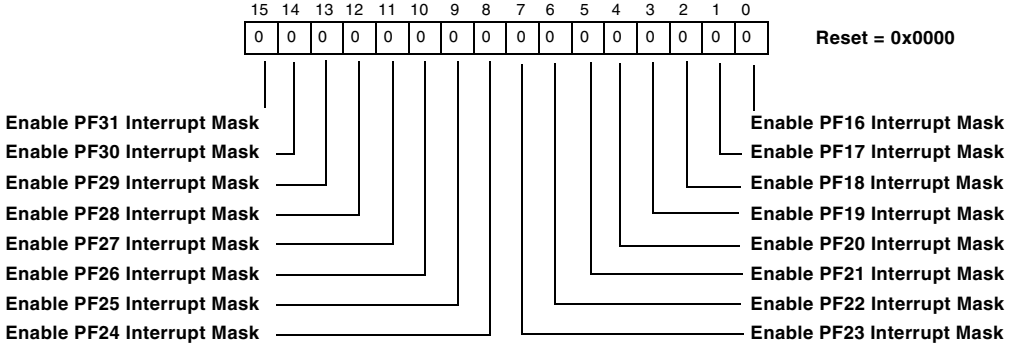

Figure 14-18. Flag (PF31 – PF16) Mask Interrupt A Data Register

Flag Mask Interrupt A Data Register (FIO2_MASKA_D)
For all bits, 1 - Enable.

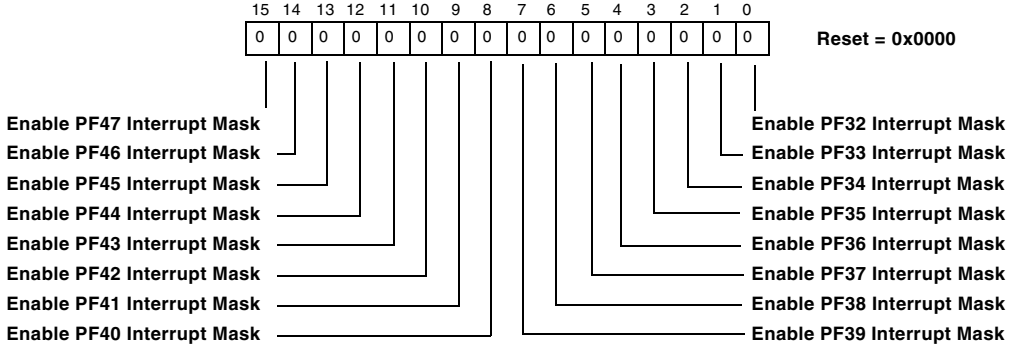

Figure 14-19. Flag (PF47 – PF32) Mask Interrupt A Data Register

Flag Mask Interrupt A Set Register (FIO0_MASKA_S)
For all bits, 1 - Set.

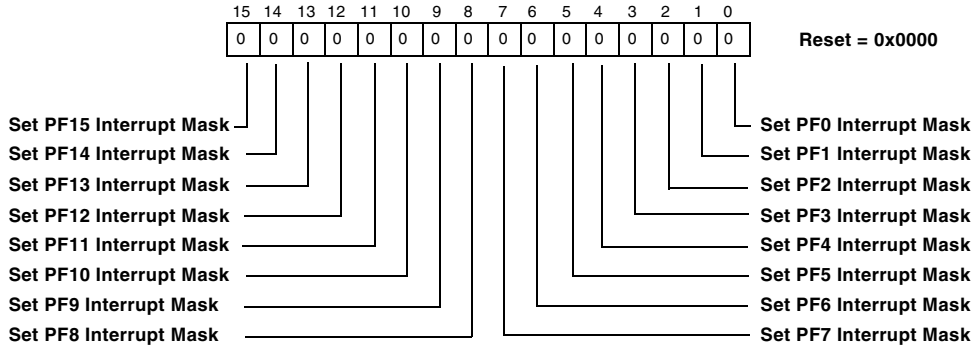

Figure 14-20. Flag (PF15 – PF0) Mask Interrupt A Set Register

Flag Mask Interrupt A Set Register (FIO1_MASKA_S)

For all bits, 1 - Set.

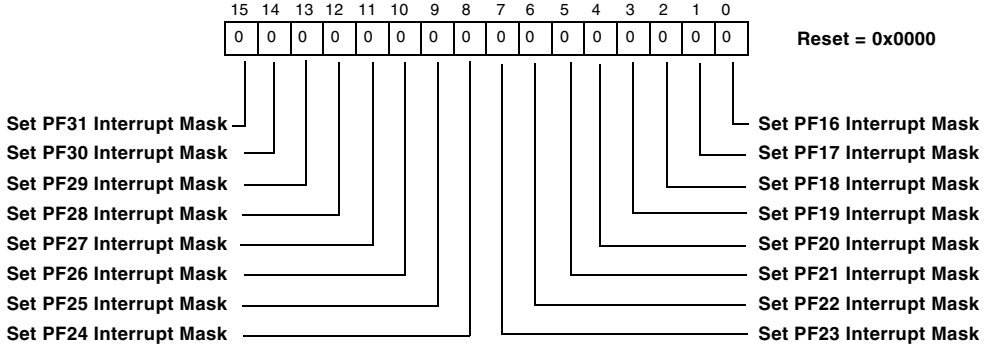

Figure 14-21. Flag (PF31 – PF16) Mask Interrupt A Set Register

Flag Mask Interrupt A Set Register (FIO2_MASKA_S)

For all bits, 1 - Set.

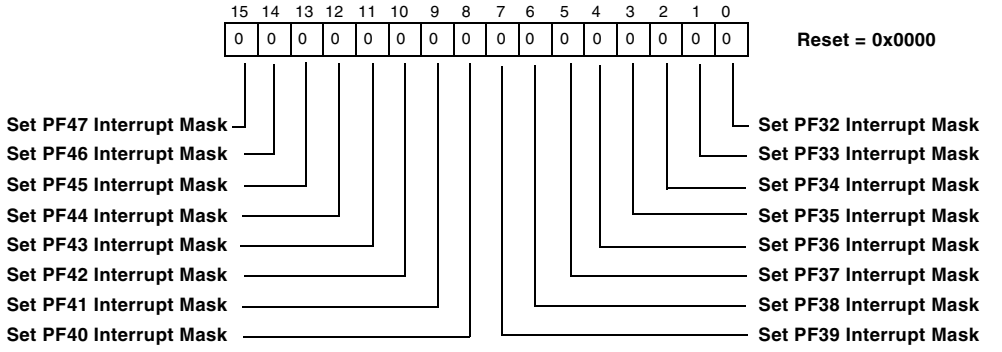

Figure 14-22. Flag (PF47 – PF32) Mask Interrupt A Set Register

Flag Mask Interrupt A Clear Register (FIO0_MASKA_C)
For all bits, 1 - Clear.

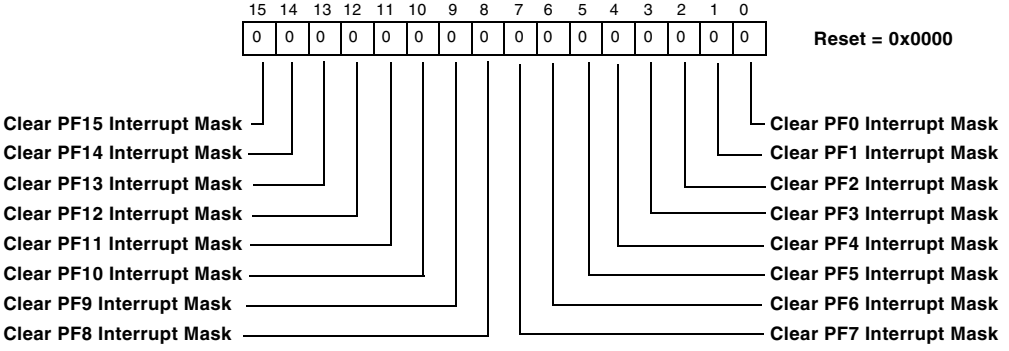

Figure 14-23. Flag (PF15 – PF0) Mask Interrupt A Clear Register

Flag Mask Interrupt A Clear Register (FIO1_MASKA_C)
For all bits, 1 - Clear.

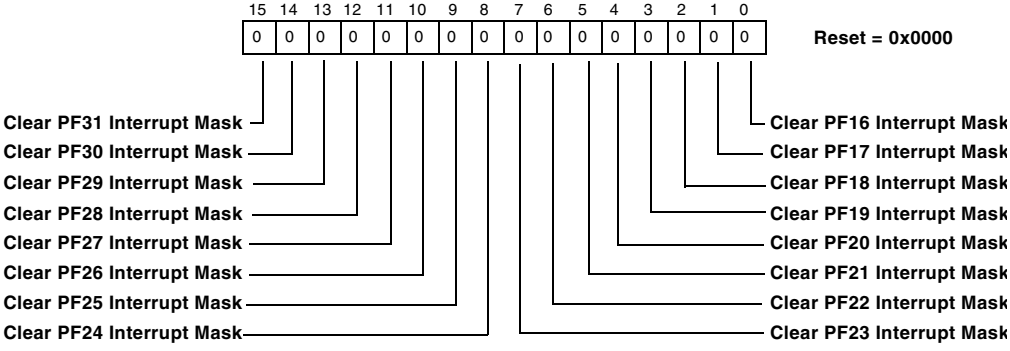

Figure 14-24. Flag (PF31 – PF16) Mask Interrupt A Clear Register

Flag Mask Interrupt A Clear Register (FIO2_MASKA_C)
For all bits, 1 - Clear.

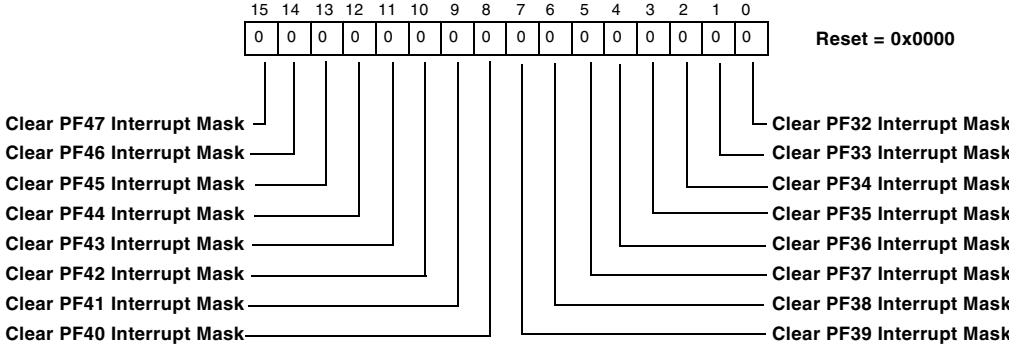

Figure 14-25. Flag (PF47 – PF32) Mask Interrupt A Clear Register

Flag Mask Interrupt A Toggle Register (FIO0_MASKA_T)
For all bits, 1 - Toggle.

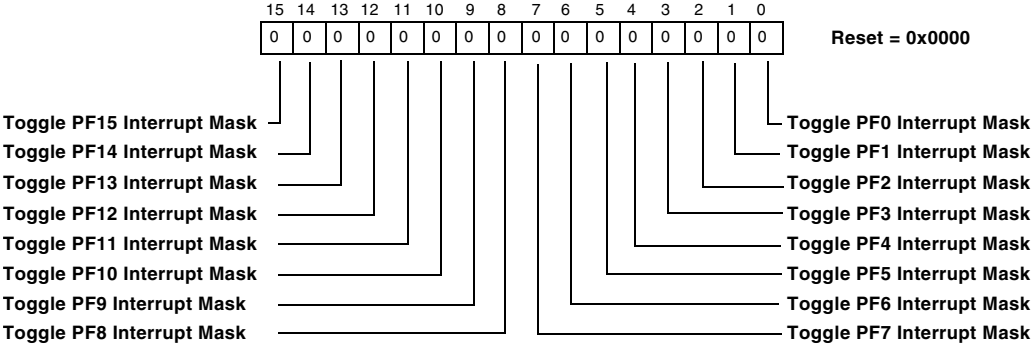

Figure 14-26. Flag (PF15 – PF0) Mask Interrupt A Toggle Register

Flag Mask Interrupt A Toggle Register (FIO1_MASKA_T)
For all bits, 1 - Toggle.

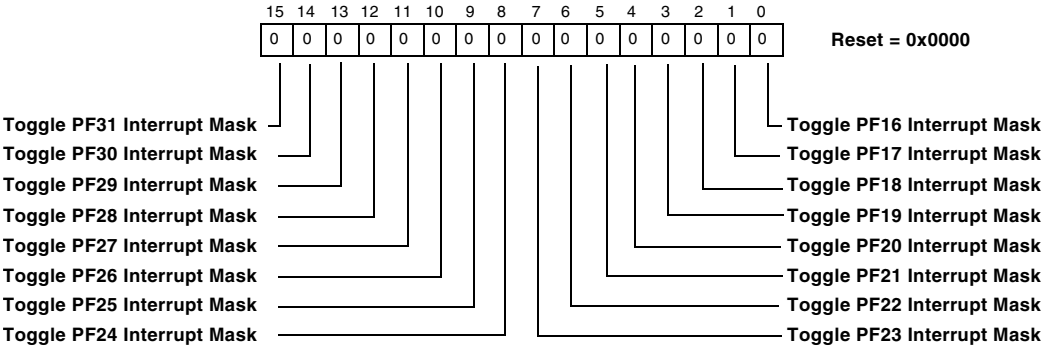

Figure 14-27. Flag (PF31 – PF16) Mask Interrupt A Toggle Register

Flag Mask Interrupt A Toggle Register (FIO2_MASKA_T)
For all bits, 1 - Toggle.

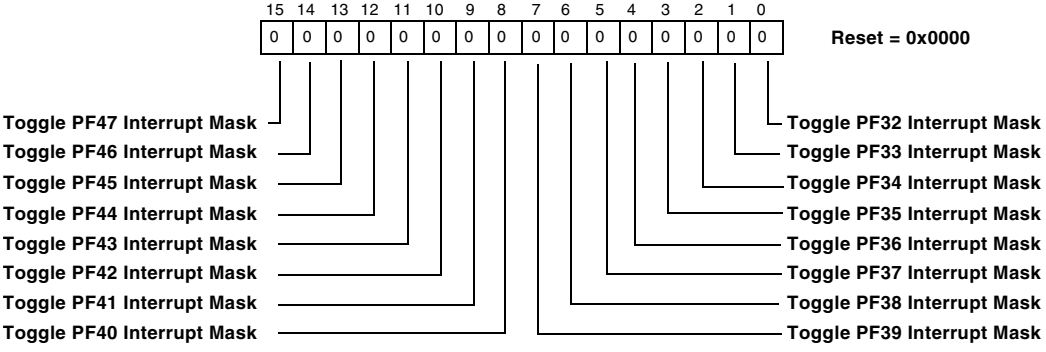

Figure 14-28. Flag (PF47 – PF32) Mask Interrupt A Toggle Register

Flag Mask Interrupt B Registers (FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)

The register sets in Figure 14-29 to Figure 14-40 on page 14-33 support Flag Interrupt B. Details can be found on page 14-19.

Flag Mask Interrupt B Data Register (FIO0_MASKB_D)
For all bits, 1 - Enable.

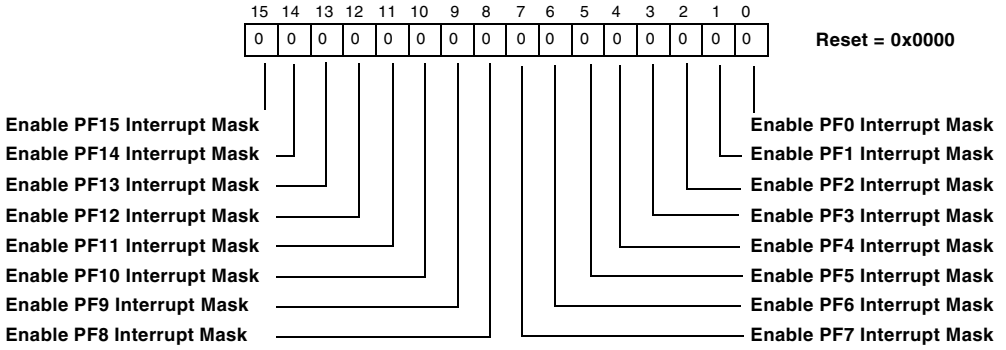

Figure 14-29. Flag (PF15 – PF0) Mask Interrupt B Data Register

Flag Mask Interrupt B Data Register (FIO1_MASKB_D)
For all bits, 1 - Enable.

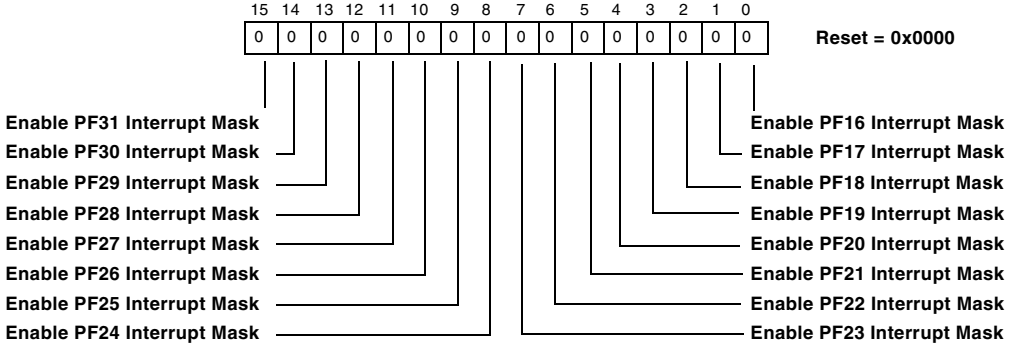

Figure 14-30. Flag (PF31 – PF16) Mask Interrupt B Data Register

Flag Mask Interrupt B Data Register (FIO2_MASKB_D)
For all bits, 1 - Enable.

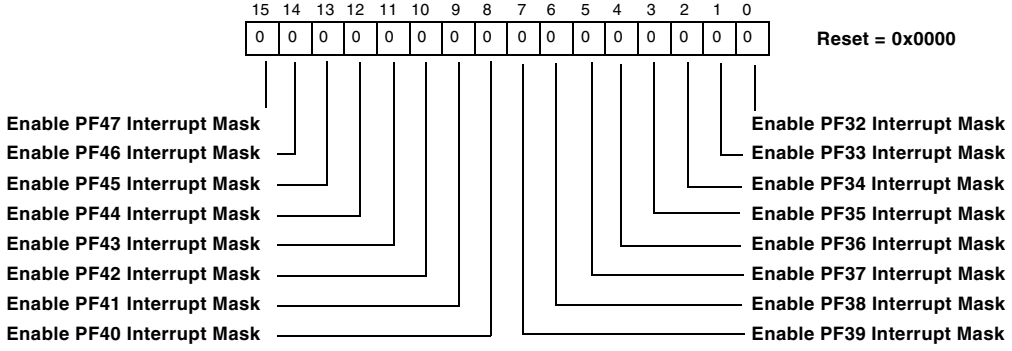

Figure 14-31. Flag (PF47 – PF32) Mask Interrupt B Data Register

Flag Mask Interrupt B Set Register (FIO0_MASKB_S)
For all bits, 1 - Set.

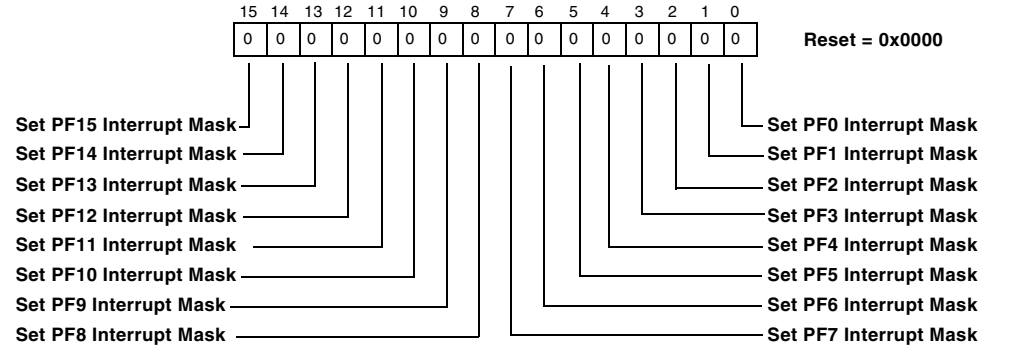

Figure 14-32. Flag (PF15 – PF0) Mask Interrupt B Set Register

Flag Mask Interrupt B Set Register (FIO1_MASKB_S)

For all bits, 1 - Set.

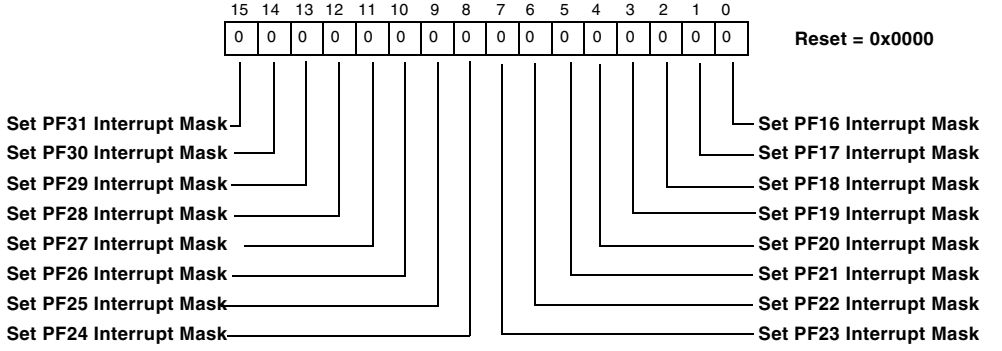

Figure 14-33. Flag (PF31 – PF16) Mask Interrupt B Set Register

Flag Mask Interrupt B Set Register (FIO2_MASKB_S)

For all bits, 1 - Set.

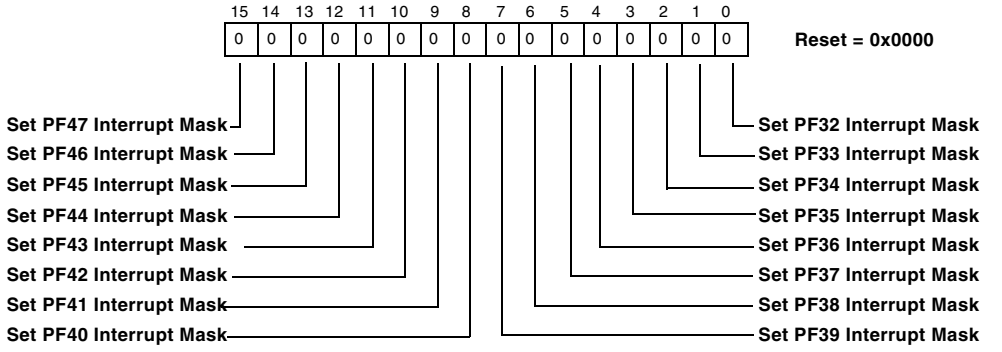

Figure 14-34. Flag (PF47 – PF32) Mask Interrupt B Set Register

Flag Mask Interrupt B Clear Register (FIO0_MASKB_C)
For all bits, 1 - Clear.

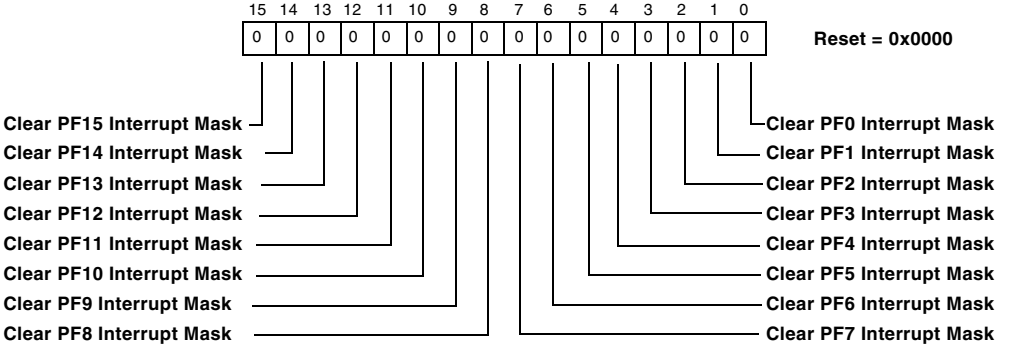

Figure 14-35. Flag (PF15 – PF0) Mask Interrupt B Clear Register

Flag Mask Interrupt B Clear Register (FIO1_MASKB_C)
For all bits, 1 - Clear.

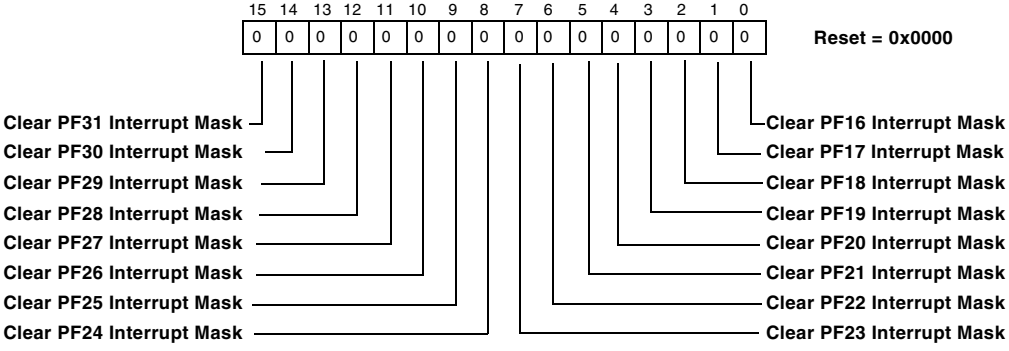

Figure 14-36. Flag (PF31 – PF16) Mask Interrupt B Clear Register

Flag Mask Interrupt B Clear Register (FIO2_MASKB_C)
For all bits, 1 - Clear.

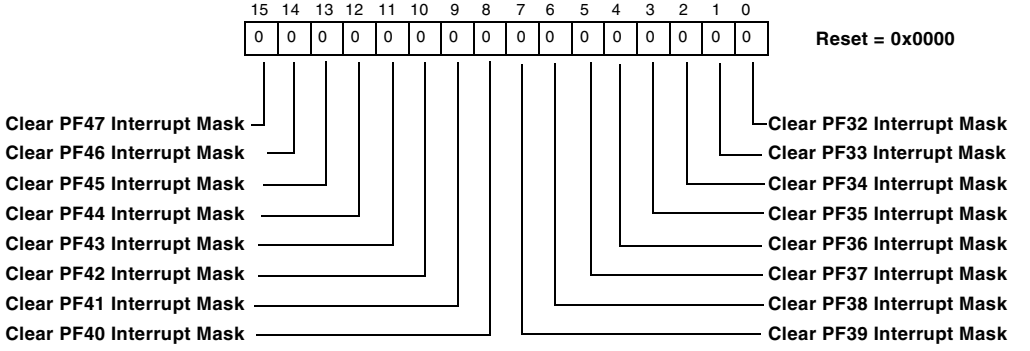

Figure 14-37. Flag (PF47 – PF32) Mask Interrupt B Clear Register

Flag Mask Interrupt B Toggle Register (FIO0_MASKB_T)
For all bits, 1 - Toggle.

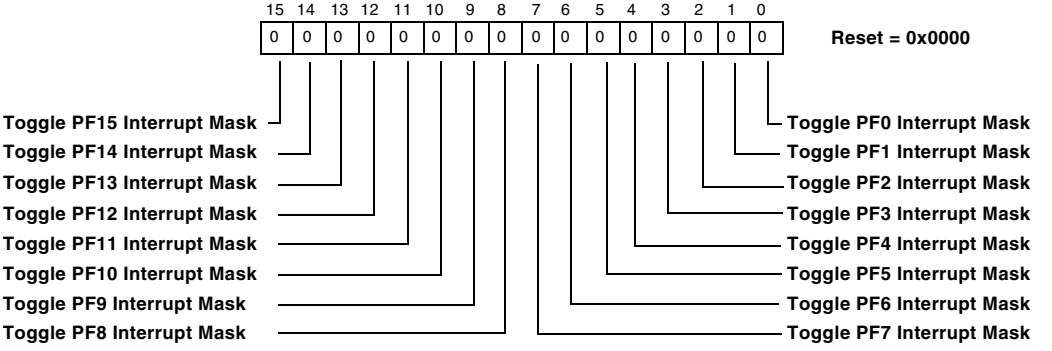

Figure 14-38. Flag (PF15 – PF0) Mask Interrupt B Toggle Register

Flag Mask Interrupt B Toggle Register (FIO1_MASKB_T)
For all bits, 1 - Toggle.

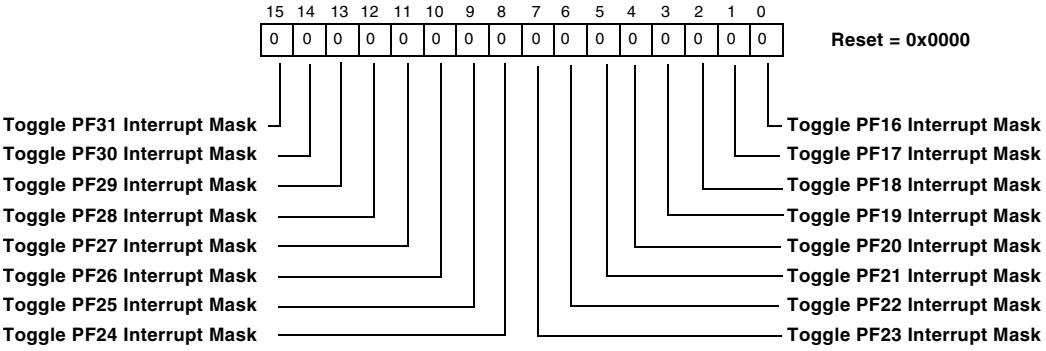

Figure 14-39. Flag (PF31 – PF16) Mask Interrupt B Toggle Register

Flag Mask Interrupt B Toggle Register (FIO2_MASKB_T)
For all bits, 1 - Toggle.

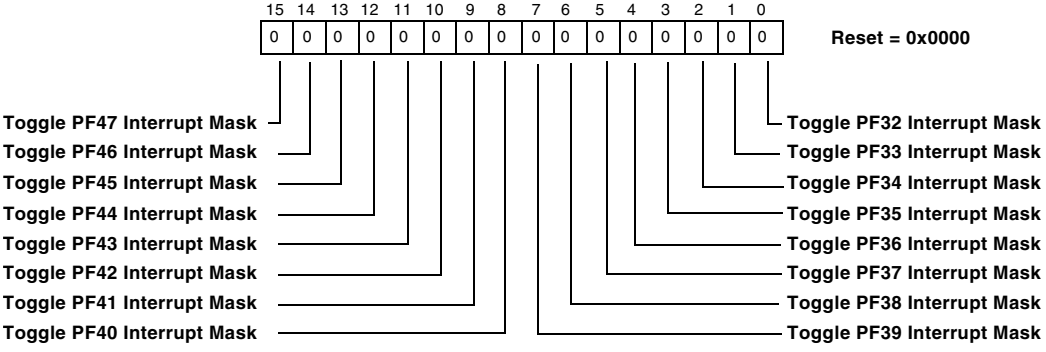

Figure 14-40. Flag (PF47 – PF32) Mask Interrupt B Toggle Register

Flag Polarity Register (FIOn_POLAR)

The Flag Polarity register is used to configure the polarity of the flag input source. To select active high or rising edge, set the bits in this register to 0.

To select active low or falling edge, set the bits in this register to 1.

This register has no effect on PFx pins that are defined as outputs. The contents of this register are cleared at reset, defaulting to active high polarity.

Flag Polarity Register (FIO0_POLAR)

For all bits, 0 - Active high or rising edge, 1 - Active low or falling edge.

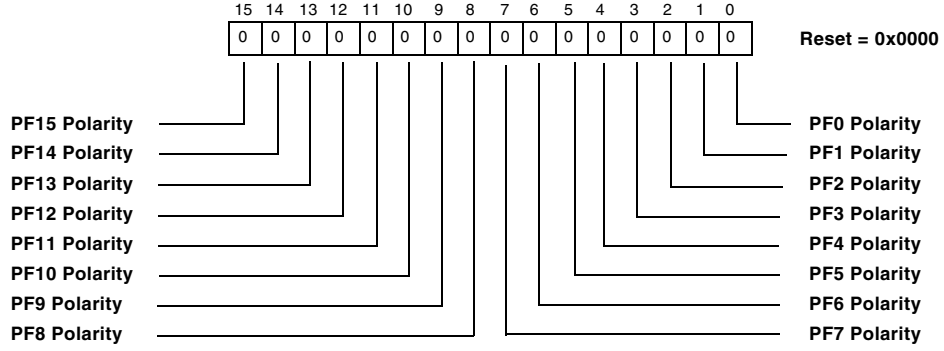

Figure 14-41. Flag (PF15 – PF0) Polarity Register

Flag Polarity Register (FIO1_POLAR)

For all bits, 0 - Active high or rising edge, 1 - Active low or falling edge.

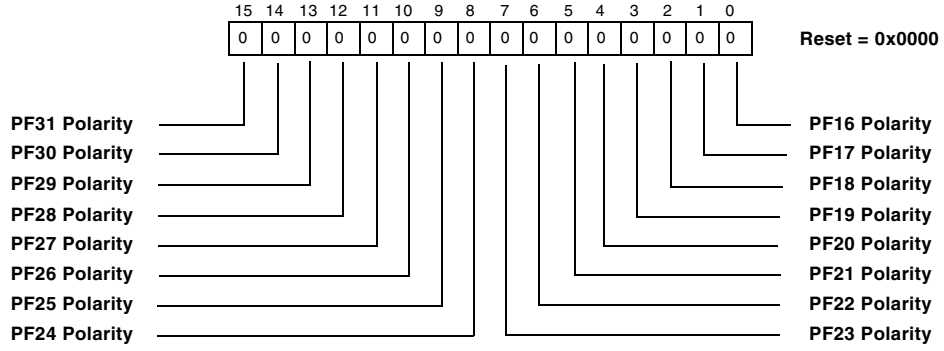

Figure 14-42. Flag (PF31 – PF16) Polarity Register

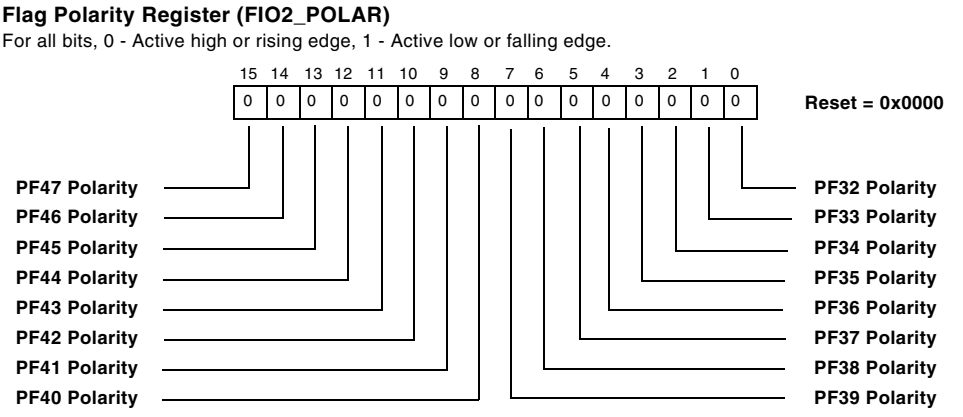

Figure 14-43. Flag (PF47 – PF32) Polarity Register

Flag Interrupt Sensitivity Register (FIOn_EDGE)

The Flag Interrupt Sensitivity register is used to configure each of the flags as either a level-sensitive or an edge-sensitive source. When using an edge-sensitive mode, an edge-detection circuit is used to prevent a situation where a short event is missed because of the system clock rate. This register has no effect on PFx pins that are defined as outputs.

The contents of this register are cleared at reset, defaulting to level sensitivity.

Flag Interrupt Sensitivity Register (FIO0_EDGE)

For all bits, 0 - Level, 1 - Edge.

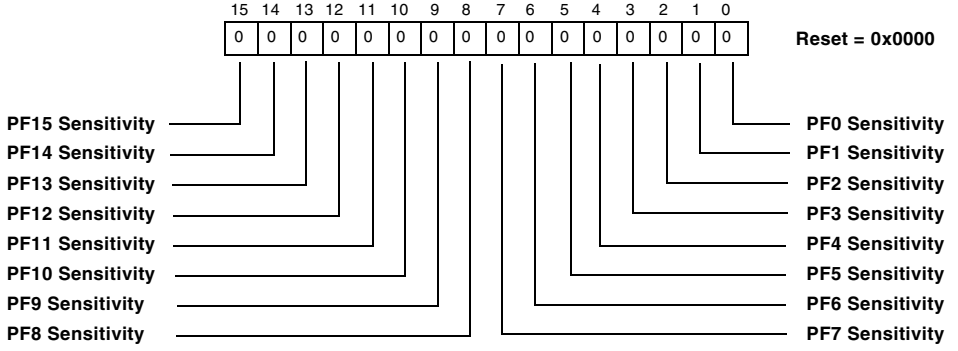

Figure 14-44. Flag (PF15 – PF0) Interrupt Sensitivity Register

Flag Interrupt Sensitivity Register (FIO1_EDGE)

For all bits, 0 - Level, 1 - Edge.

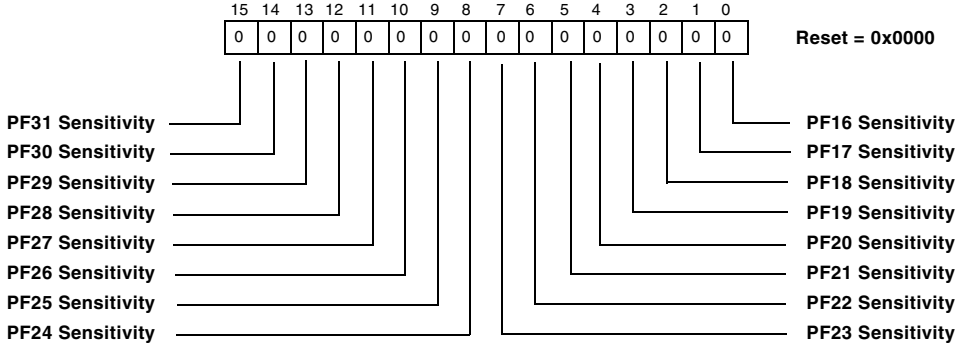

Figure 14-45. Flag (PF31 – PF16) Interrupt Sensitivity Register

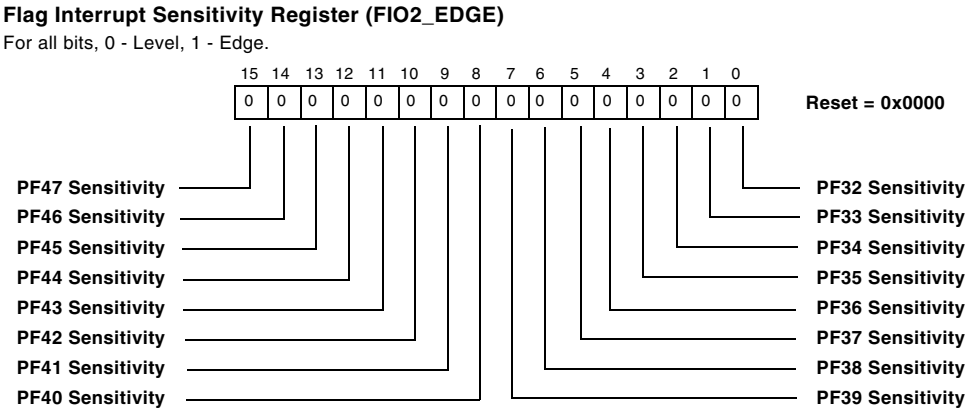

Figure 14-46. Flag (PF47 – PF32) Interrupt Sensitivity Register

Flag Set on Both Edges Register (FIOn_BOTH)

The Flag Set on Both Edges register is used to enable interrupt generation on both rising and falling edges.

When a given PFx pin has been set to edge-sensitive in the Flag Interrupt Sensitivity register, setting the PFx pin's bit in the Flag Set on Both Edges register to Both Edges results in an interrupt being generated on both the rising and falling edges. This register has no effect on PFx pins that are defined as level-sensitive or as outputs.

The contents of this register are cleared at reset, defaulting to single-edge sensitivity.

Flag Set on Both Edges Register (FIO0_BOTH)

For all bits, when enabled for edge sensitivity, 0 - Single edge, 1 - Both edges.

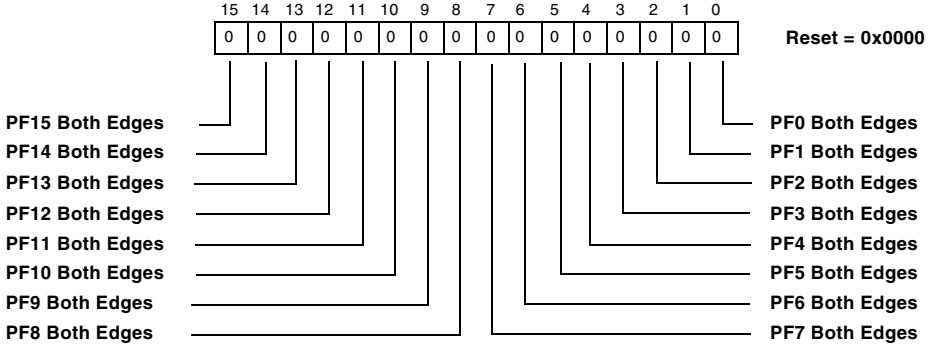

Figure 14-47. Flag (PF15 – PF0) Set on Both Edges Register

Flag Set on Both Edges Register (FIO1_BOTH)

For all bits, when enabled for edge sensitivity, 0 - Single edge, 1 - Both edges.

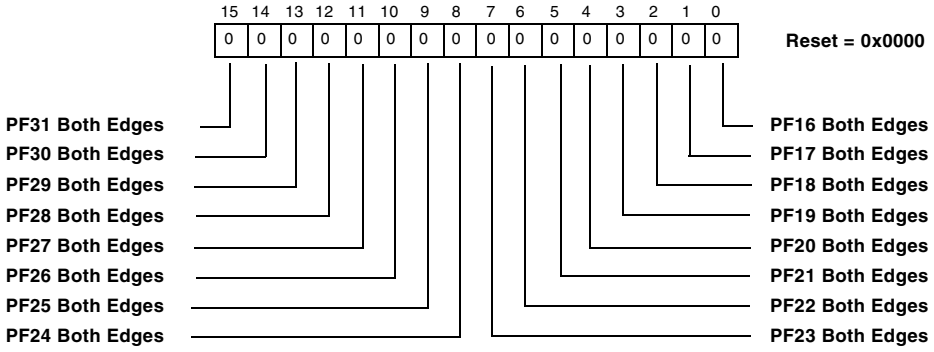

Figure 14-48. Flag (PF31 – PF16) Set on Both Edges Register

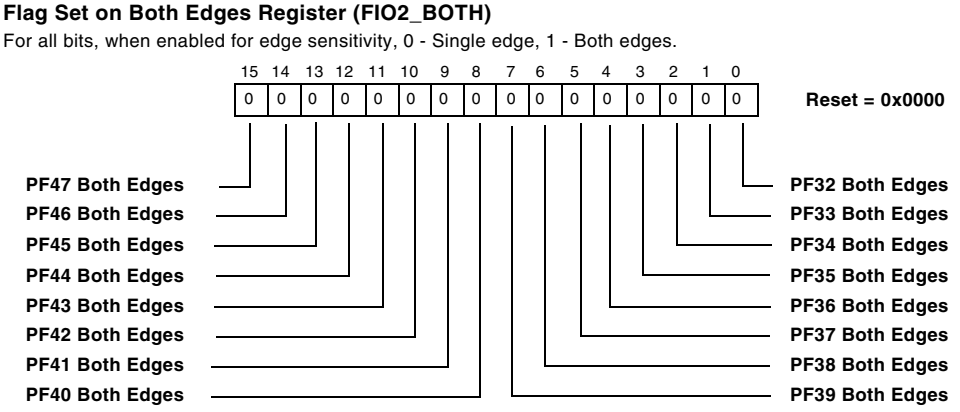

Figure 14-49. Flag (PF47 – PF32) Set on Both Edges Register

Flag Input Enable Register (FIOn_INEN)

The Flag Input Enable register is used to enable the input buffers on any flag pin that is being used as an input. Leaving the input buffer disabled eliminates the need for pullups and pulldowns when a particular PFx pin is not used in the system. By default, the input buffers are disabled.

Flag Input Enable Register (FIO0_INEN)
For all bits, 0 - Input Buffer Disabled, 1 - Input Buffer Enabled.

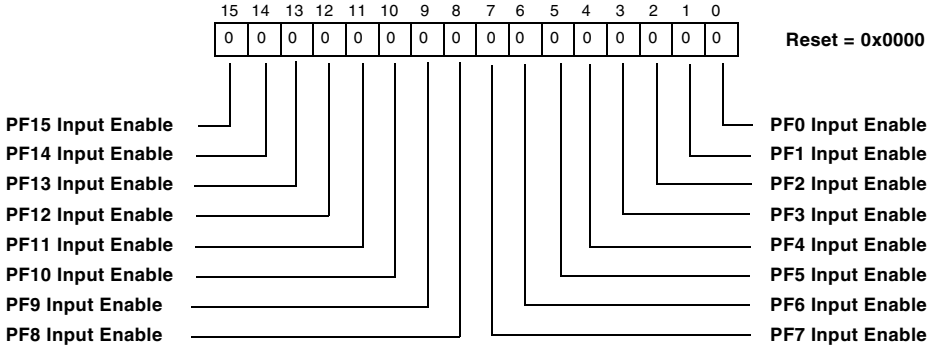

Figure 14-50. Flag (PF15 – PF0) Input Enable Register

Flag Input Enable Register (FIO1_INEN)
For all bits, 0 - Input Buffer Disabled, 1 - Input Buffer Enabled.

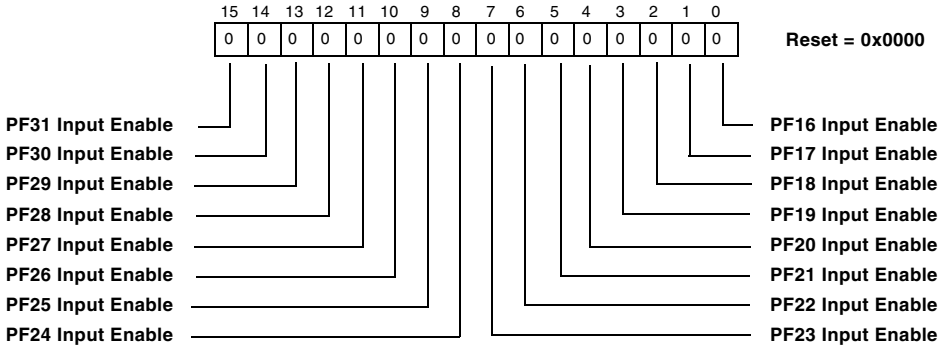

Figure 14-51. Flag (PF31 – PF16) Input Enable Register

Flag Input Enable Register (FIO2_INEN)
For all bits, 0 - Input Buffer Disabled, 1 - Input Buffer Enabled.

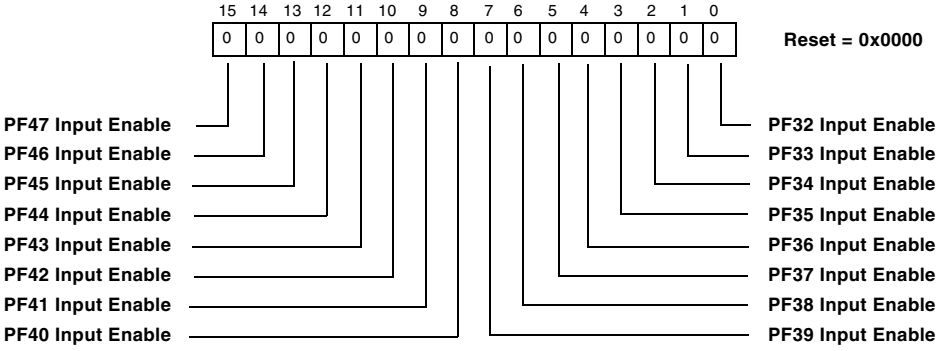

Figure 14-52. Flag (PF47 – PF32) Input Enable Register

Note that if the PFx pin is being used as an input, the corresponding bit in the Flag Input Enable register must be set. Otherwise, changes in the external flag pins will not be recognized by the processor.

Performance/Throughput

The PFx pins are synchronized to the system clock (SCLK). When configured as outputs, the programmable flags can transition once every system clock cycle.

When configured as inputs, the overall system design should take into account the potential latency between the core and system clocks. Changes in the state of PFx pins have a latency of 3 SCLK cycles before being detectable by the processor. When configured for level-sensitive interrupt generation, there is a minimum latency of 4 SCLK cycles between the time the flag is asserted and the time that program flow is interrupted. When configured for edge-sensitive interrupt generation, an additional SCLK

cycle of latency is introduced, giving a total latency of 5 SCLK cycles between the time the edge is asserted and the time that the core program flow is interrupted.

15 TIMERS

The processor features twelve identical 32-bit general-purpose timers, and each core also has a core timer, and a watchdog timer.

The general-purpose timers can be individually configured in any of three modes:

- Pulse Width Modulation (PWM_OUT) mode
- Pulse Width Count and Capture (WDTH_CAP) mode
- External Event (EXT_CLK) mode

The core timer is available for each core and can be used to generate periodic interrupts for a variety of system timing functions.

A dedicated watchdog timer for each core can be used to implement a software watchdog function. A software watchdog can improve system availability by generating an event to the Blackfin core if the timer expires before being updated by software.

General-Purpose Timers

Each general-purpose timer has one dedicated bidirectional chip pin, TMRx. This pin functions as an output pin in the PWM_OUT mode and as an input pin in the WDTH_CAP and EXT_CLK modes. To provide these functions, each timer has four registers. For range and precision, the Timer Counter (TIMERx_COUNTER), Timer Period (TIMERx_PERIOD), and Timer Pulse Width (TIMERx_WIDTH) registers are 32 bits wide, as shown in Figure 15-1.

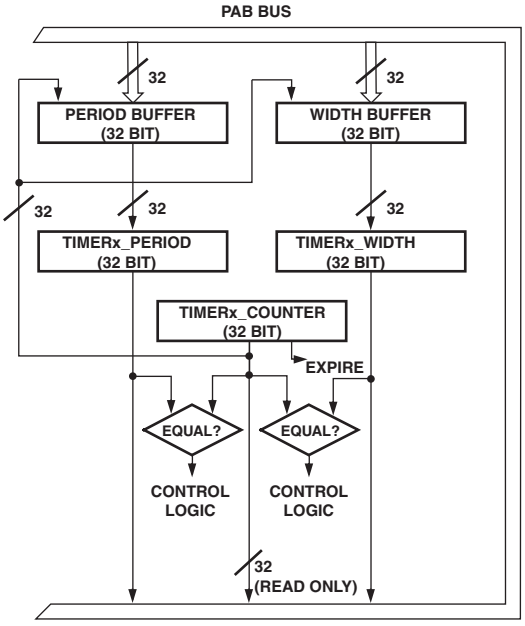

Figure 15-1. Timer Block Diagram

The registers for each general-purpose timer are:

- Timer Configuration (TIMERx_CONFIG) registers
- Timer Counter (TIMERx_COUNTER) registers
- Timer Period (TIMERx_PERIOD) registers
- Timer Pulse Width (TIMERx_WIDTH) registers

When clocked internally, the clock source is the processor's peripheral clock (SCLK). Assuming the peripheral clock is running at 133 MHz, the maximum period for the timer count is $((2^{32}-1) / 133\ \text{MHz}) = 32.2$ seconds.

There are two separate pairs of Timer Enable and Disable registers. TMRS8_ENABLE and TMRS8_DISABLE control the first eight timers, which are

numbered 0 to 7. TMRS4_ENABLE and TMRS4_DISABLE control the remaining four timers with numbers 8 to 11. The Timer Enable registers contain 8 and 4 “write-one-to-set” control bits for enabling their corresponding timers. Similarly, the Timer Disable registers contain 8 and 4 “write-one-to-clear” bits for disabling their corresponding timers. Either the Timer Enable or the Timer Disable registers can be read back to check the enable status of the timers. A 1 indicates that the corresponding timer is enabled. The timer starts counting three SCLK cycles after the TIMENx bit is set.

The Timer Status (TMRS8_STATUS and TMRS4_STATUS) registers contain an Interrupt Latch bit (TIMILx) and an Overflow/Error Indicator bit (TOVF_ERRx) for each timer. These sticky bits are set by the timer hardware and may be polled by software. They need to be cleared by software explicitly, by writing a 1 to the bit.

To enable a timer’s interrupts, set the IRQ_ENA bit in the timer’s Configuration (TIMERx_CONFIG) register and unmask the timer’s interrupt by setting the corresponding bits of the IMASK and SICx_IMASK registers. With the IRQ_ENA bit cleared, the timer does not set its Interrupt Latch (TIMILx) bits. To poll the TIMILx bits without permitting a timer interrupt, programs can set the IRQ_ENA bit while leaving the timer’s interrupt masked.

With interrupts enabled, make sure that the interrupt service routine (ISR) clears the TIMILx latch before the RTI instruction, to ensure that the interrupt is not reissued. To make sure that no timer event is missed, the latch should be reset at the very beginning of the interrupt routine when in External Clock (EXT_CLK) mode. To enable timer interrupts, set the IRQ_ENA bit in the proper Timer Configuration (TIMERx_CONFIG) register.

For each Parallel Peripheral Interface (PPI) that is enabled for internally generated frame syncs and CLK_SEL = 1, up to two timers are configured for PWM_OUT mode and used to generate the needed frame sync signals. For more information, refer to Chapter 11, “Parallel Peripheral Interface”.

To enter the mode where the PPI and timers are used together, first program the PPI, then write the Timer Configuration register for the timer or timers you want to use. Enable the PPI and then enable the timers. When done, disable the PPI and then disable the timers. To exit this mode, reprogram the PPI and then write the Timer Configuration registers to 0.

Timer Registers

The timer peripheral module provides general-purpose timer functionality. It consists of twelve identical timer units.

The twelve timers are arranged into two groups, Group 1 and Group 2. Group 1 consists of eight timers, and group 2 consists of four timers. Each group has the following registers:

- `TIMERx_CONFIG[15:0]` – Timer Configuration register
- `TIMERx_WIDTH[31:0]` – Timer Pulse Width register
- `TIMERx_PERIOD[31:0]` – Timer Period register
- `TIMERx_COUNTER[31:0]` – Timer Counter register

Three registers in each group are shared among these timers, where y = 8: Group 1, y = 4: Group 2:

- `TMRSy_ENABLE[15:0]` – Timer Enable register
- `TMRSy_DISABLE[15:0]` – Timer Disable register
- `TMRSy_STATUS[31:0]` – Timer Status register

The size of accesses is enforced. A 32-bit access to a Timer Configuration register or a 16-bit access to a Timer Pulse Width, Timer Period, or Timer Counter register results in an MMR error. Both 16-bit and 32-bit accesses are allowed for the Timer Enable, Timer Disable, and Timer Status registers. On a 32-bit read, the upper word returns all zeros.

Timer Enable (TMRSy_ENABLE) Registers

The Timer Enable registers (Figure 15-2) allow 8 or 4 timers, for y = 8 or 4 respectively, to be enabled simultaneously in order to make them run synchronously. For each timer there is a single W1S control bit. Writing a 1 enables the corresponding timer, writing a 0 has no effect. The twelve bits can be set individually or in any combination. A read of the Timer Enable registers shows the status of the enable for the corresponding timer. A 1 indicates that the timer is enabled. All unused bits return 0 when read.

Timer Disable (TMRSy_DISABLE) Registers

The Timer Disable registers (Figure 15-3) allow 8 or 4 timers, for y = 8 or 4 respectively, to be disabled simultaneously. For each timer there is a single W1C control bit. Writing a 1 disables the corresponding timer; writing a 0 has no effect. The twelve bits can be cleared individually or in any combination. A read of the Timer Disable registers show the status of the disable for the corresponding timer. A "0" indicates that the timer is disabled. All unused bits return 0 when read.

In `PWM_OUT` mode, a write of a 1 to `TMRSy_DISABLE` does not stop the corresponding timer immediately. Rather, the timer continues running and stops cleanly at the end of the current period (if `PERIOD_CNT = 1`) or pulse (if `PERIOD_CNT = 0`). If necessary, the processor can force a timer in `PWM_OUT` mode to stop immediately by first writing a 1 to the corresponding bit in `TMRSy_DISABLE`, and then writing a 1 to the corresponding `TRUNx` bit in `TMRSy_STATUS`. See "Stopping the Timer in PWM_OUT Mode" on page 15-23.

In `WDTH_CAP` and `EXT_CLK` modes, a write of a 1 to `TMRSy_DISABLE` stops the corresponding timer immediately.

Timer Enable Registers (TMRSy_ENABLE)

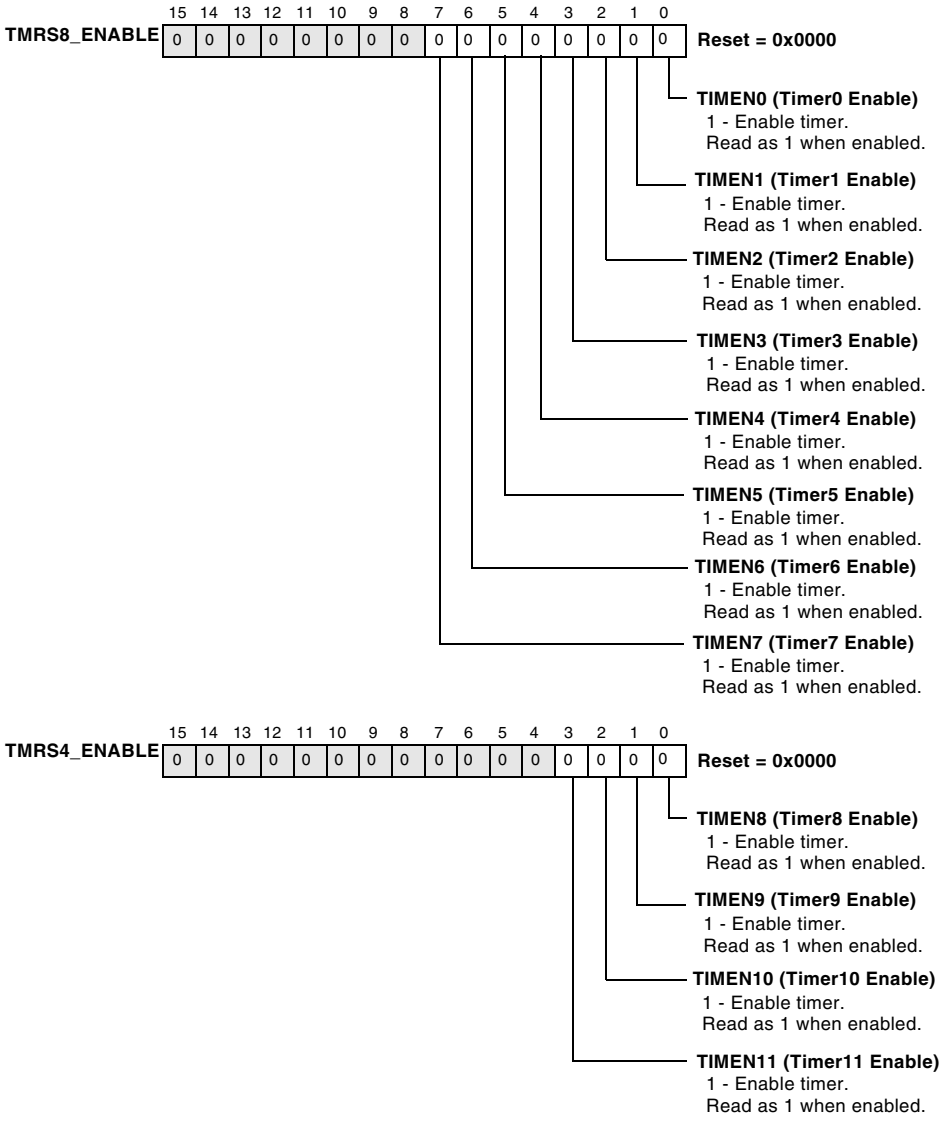

Figure 15-2. Timer Enable Registers

Timer Disable Registers (TMRSy_DISABLE)

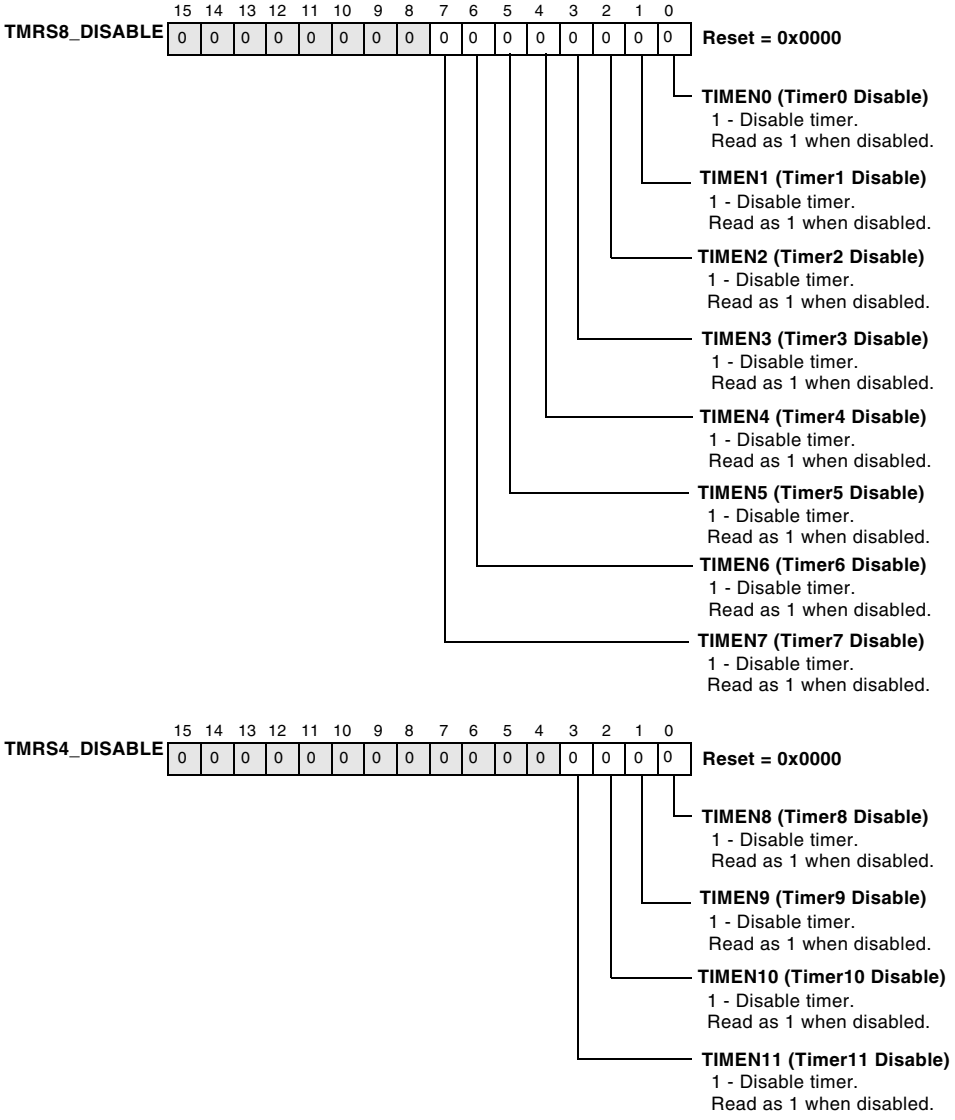

Figure 15-3. Timer Disable Registers

Timer Status (TMRSy_STATUS) Registers

The two Timer Status registers (Figure 15-4 and Figure 15-5) indicate the status of the 8 or 4 timers, for y = 8 or 4 respectively, and can be used to check the status of all timers in the group with a single read. Status bits are sticky and W1C. The TRUN bits can clear themselves, which they do when a PWM_OUT mode timer stops at the end of a period. During a Status Register read access, all reserved or unused bits return a 0.

Each Timer generates a unique processor Interrupt Request signal, which is gated by the corresponding IRQ_ENA bit in the TIMERx_CONFIG register. The shared Status Registers (TMRSy_STATUS) latch these interrupts so that the user can determine the interrupt source without reference to the unique interrupt signal (for example, in the case where all twelve timers have been assigned to the same interrupt priority). Interrupt bits are sticky and must be cleared by the interrupt service routine to assure that the interrupt is not re-issued.

The TIMILx bits work along with the IRQ_ENA bit of the Timer Configuration register to indicate interrupt requests. If an interrupt condition or error occurs and IRQ_ENA is set, then the TIMILx bit is set and the interrupt to the core is asserted. This interrupt may be masked by the system interrupt controller. If an interrupt condition or error occurs and IRQ_ENA is cleared, then the TIMILx bit is not set and the interrupt is not asserted. If TIMILx is already set and IRQ_ENA is written to 0, TIMILx stays set and the interrupt stays asserted. See Figure 15-26 on page 15-43.

The read value of the TRUNx bits reflects the timer slave enable status in all modes: TRUNx set means running and TRUNx cleared means stopped. While the TIMENx or TIMDISx bits in the TMRSy_ENABLE and TMRSy_DISABLE registers reflect whether a timer is enabled, the TRUNx bits indicate whether the timer is actually running. In WDTH_CAP and EXT_CLK modes, reads from TIMENx and TRUNx always return the same value.

A W1C operation to the TMRSy_DISABLE register disables the corresponding timer in all modes. In PWM_OUT mode, a disabled timer continues running until the ongoing period (PERIOD_CNT = 1) or pulse (PERIOD_CNT = 0) completes. During this final period the TIMENx bit returns 0, but the TRUNx bit still reads as a 1. See Figure 15-12 on page 15-19. In this state only, TRUNx becomes a W1C bit. During this final period with the timer disabled, writing a 1 to TRUNx clears TRUNx and stops the timer immediately without waiting for the Timer Counter to reach the end of its current cycle.

Writing the TRUNx bits has no effect in other modes or when a timer has not been enabled. Writing the TRUNx bits to 1 in PWM_OUT mode has no effect on a timer that has not first been disabled.

Timer Status Register (TMRS8_STATUS)

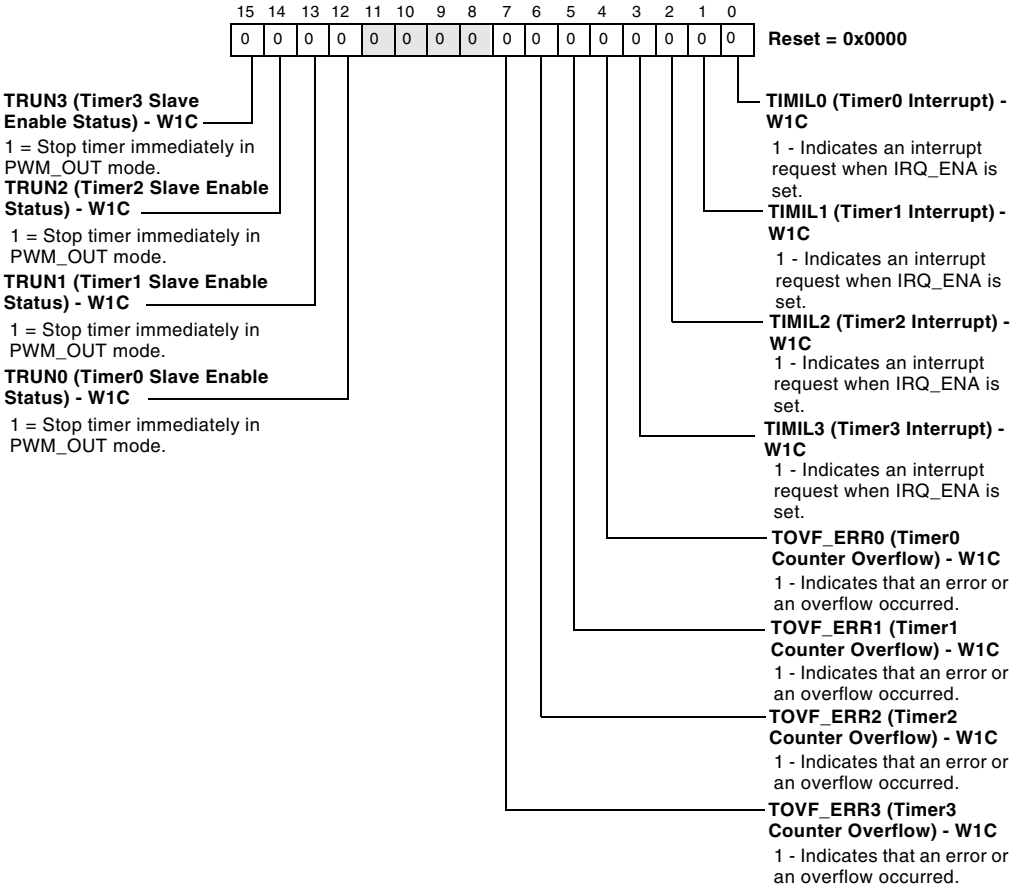

Figure 15-4. Timer Group 1 Low 16-Bit Status Register

Timer Status Register (TMRS8_STATUS)

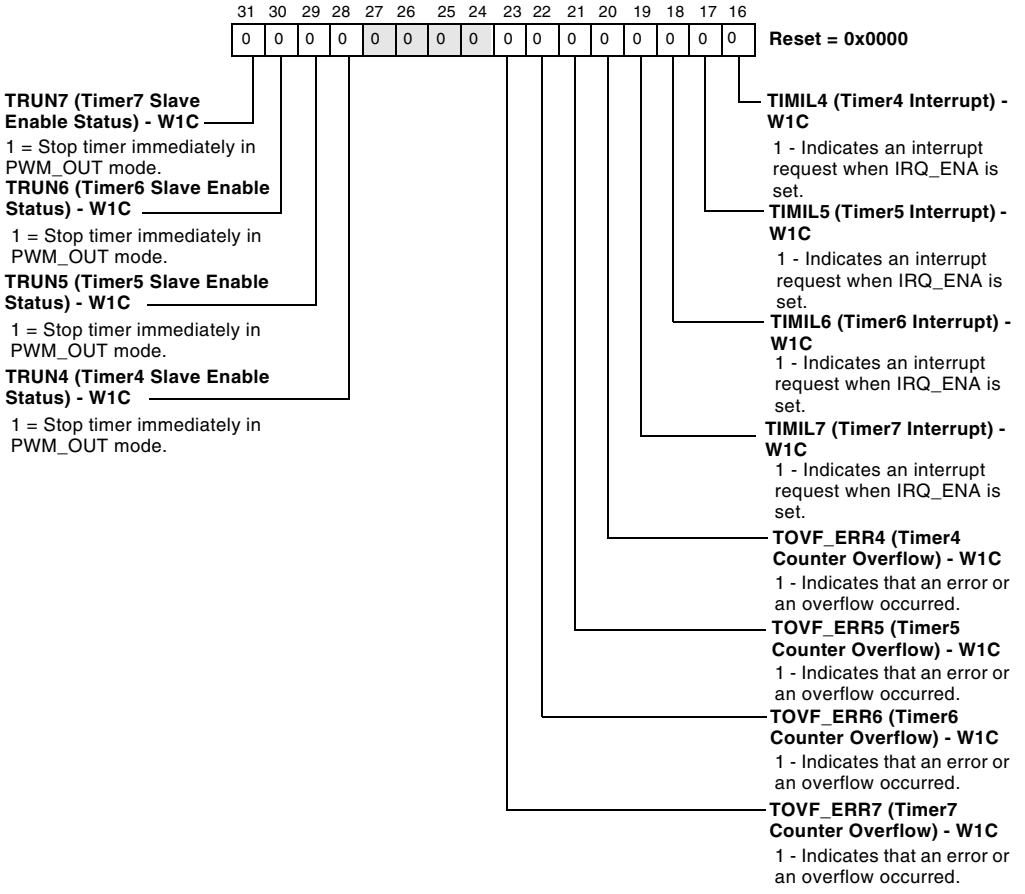

Figure 15-5. Timer Group 1 High 16-Bit Status Register

Timer Status Register (TMRS4_STATUS)

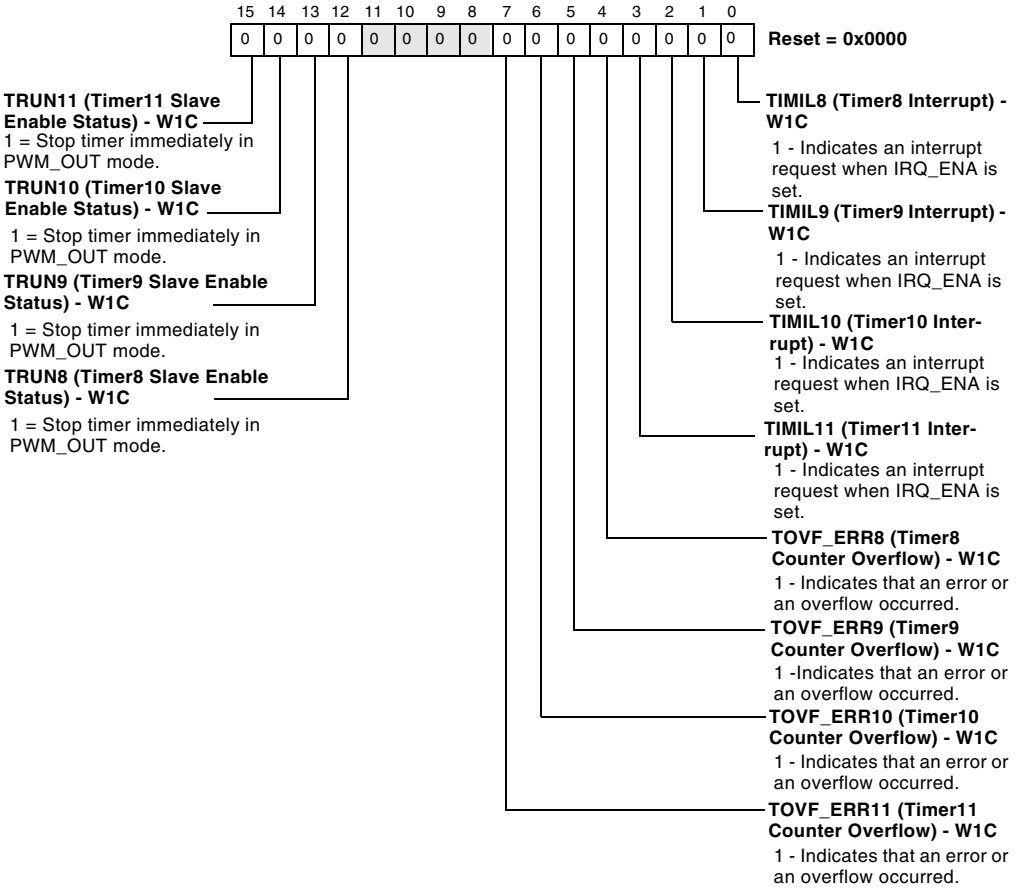

Figure 15-6. Timer Group 2 Status Register

Timer Configuration (TIMERx_CONFIG) Registers

The operating mode for each timer is specified by its configuration register. The `TIMERx_CONFIG` register may be written only when the timer is not running. After disabling the timer in `PWM_OUT` mode, make sure the timer

has stopped running by checking its TRUN bit in TMRSy_STATUS before attempting to reprogram TIMERx_CONFIG. The TIMERx_CONFIG registers may be read at any time. The ERR_TYP field is read-only. It is cleared at reset and when the timer is enabled. Each time TOVF_ERRx is set, ERR_TYP[1:0] is loaded with a code that identifies the type of error that was detected. This value is held until the next error or timer enable occurs. For an overview of error conditions, see Table 15-1 on page 15-45. Figure 15-7 provides bit descriptions. The TIMERx_CONFIG register also controls the behavior of the TMRx pin, which becomes an output in PWM_OUT mode (TMODE = 01) when the OUT_DIS bit is cleared.

Timer Configuration Registers (TIMERx_CONFIG)

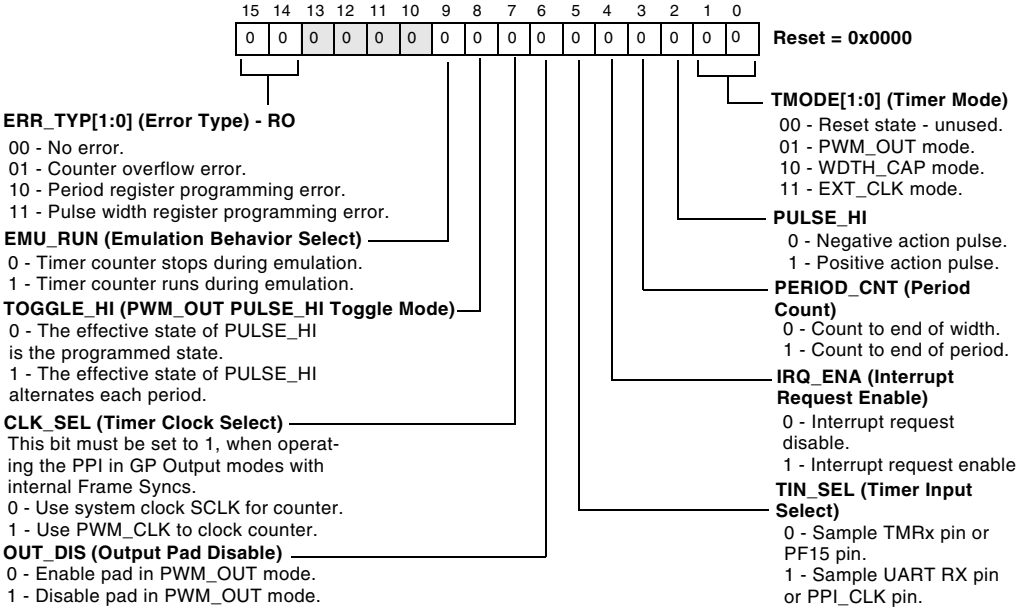

Figure 15-7. Timer Configuration Registers

Timer Counter (TIMERx_COUNTER) Registers

These read-only registers retain their state when disabled. When enabled, the Timer Counter register (Figure 15-8) is re-initialized by hardware based on configuration and mode.

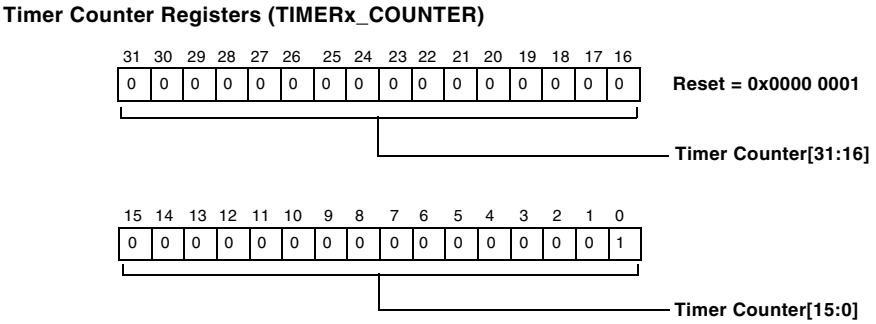

Figure 15-8. Timer Counter Registers

The Timer Counter register may be read at any time (whether the timer is running or stopped), and it returns a coherent 32-bit value. Depending on the operation mode, the incrementing counter can be clocked by four different sources: SCLK, the TMRx pin, the programmable flag pin PF15, or the parallel peripheral clock PPI_CLK.

While the processor core is being accessed by an external emulator debugger, all code execution stops. By default, the TIMERx_COUNTER also halts its counting during an emulation access in order to remain synchronized with the DSP software. While stopped, the count does not advance: in PWM_OUT mode, the TMRx pin waveform is “stretched”; in WDTH_CAP mode, measured values are incorrect; in EXT_CLK mode, input events on TMRx may be missed. All other Timer functions such as register reads and writes, interrupts previously asserted (unless cleared), and the loading of TIMERx_PERIOD and TIMERx_WIDTH in WDTH_CAP mode remain active during an emulation stop.

Some applications may require the timer to continue counting asynchronously to the emulation-halted processor core. Set the EMU_RUN bit in TIMERx_CONFIG to enable this behavior.

Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers

When a timer is enabled and running, and the software writes new values to the Timer Period register and the Timer Pulse Width register, the writes are buffered and do not update the registers until the end of the current period (when the Timer Counter register equals the Timer Period register).

Usage of the Timer Period and Timer Pulse Width registers (Figure 15-9 and Figure 15-10) varies depending on the mode of the timer:

- In Pulse Width Modulation mode (PWM_OUT), both the Timer Period and Timer Pulse Width register values can be updated “on the fly” since the Timer Period and Timer Pulse Width (duty cycle) register values change simultaneously.
- In Pulse Width and Period Capture mode (WDTH_CAP), the Timer Period and Timer Pulse Width values are captured at the appropriate time. The Timer Period and Timer Pulse Width registers are then updated simultaneously from their respective buffers. Both registers are read-only in this mode.
- In External Event Capture mode (EXT_CLK), the Timer Period register is writable and can be updated “on the fly”. The Timer Pulse Width register is not used.

If new values are not written to the Timer Period register or the Timer Pulse Width register, the value from the previous period is re-used. Writes to the 32-bit Timer Period register and Timer Pulse Width register are atomic; it is not possible for the high word to be written without the low word also being written.

Values written to the Timer Period registers or Timer Pulse Width registers are always stored in the buffer registers. Reads from the Timer Period or Timer Pulse Width registers always return the current, active value of period or pulse width. Written values are not read back until they become active. When the timer is enabled, this is after the Timer Period and Timer Pulse Width registers are updated from their respective buffers at the end of the current period. See Figure 15-1 on page 15-2.

When the timer is disabled, writes to the buffer registers are immediately copied through to the Timer Period or Timer Pulse Width register so that they will be ready for use in the first timer period. For example, to change the values for the Timer Period and/or Timer Pulse Width registers in order to use a different setting for each of the first three timer periods after the timer is enabled, the procedure to follow is:

1. Program the first set of register values.
2. Enable the timer.
3. Immediately program the second set of register values.
4. Wait for the first timer interrupt.
5. Program the third set of register values.

Each new setting is then programmed when a timer interrupt is received.

In `PWM_OUT` mode with very small periods (less than 10 counts), there may not be enough time between updates from the buffer registers to write both the Timer Period register and the Timer Pulse Width register. The next period may use one old value and one new value. In order to prevent Pulse Width >= Period errors, write the Timer Pulse Width register before the Timer Period register when decreasing the values, and write the Timer Period register before the Timer Pulse Width register when increasing the value.

Timer Period Registers (TIMERx_PERIOD)

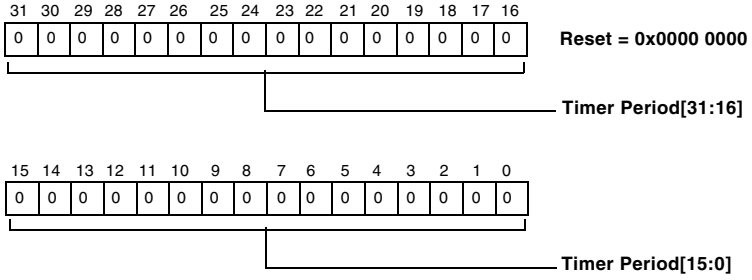

Figure 15-9. Timer Period Registers

Timer Width Registers (TIMERx_WIDTH)

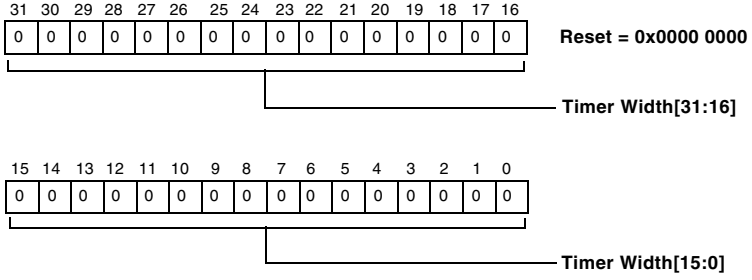

Figure 15-10. Timer Width Registers

Using the Timer

To enable an individual timer, set that timer's `TIMEN` bit in the `TMRSy_ENABLE` register. To disable an individual timer, set that timer's `TIMDIS` bit in the `TMRSy_DISABLE` register.

Before enabling a timer, always program the corresponding Timer Configuration (`TIMERx_CONFIG`) register. This register defines the timer operating

mode, the polarity of the TMRx pin, and the timer interrupt behavior. Do not alter the operating mode while the timer is running.

Examples of timer enable and disable timing appear in Figure 15-11, Figure 15-12, and Figure 15-13.

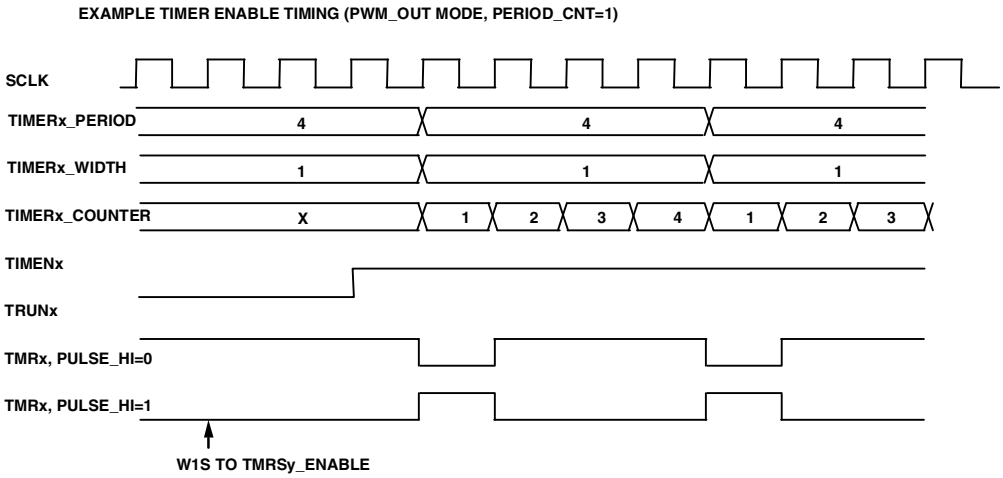

Figure 15-11. Timer Enable Timing

When timers are disabled, the Timer Counter registers retain their state; when a timer is re-enabled, the Timer Counter is re-initialized based on the operating mode. The Timer Counter registers are read-only. Software cannot overwrite or preset the Timer Counter value directly.

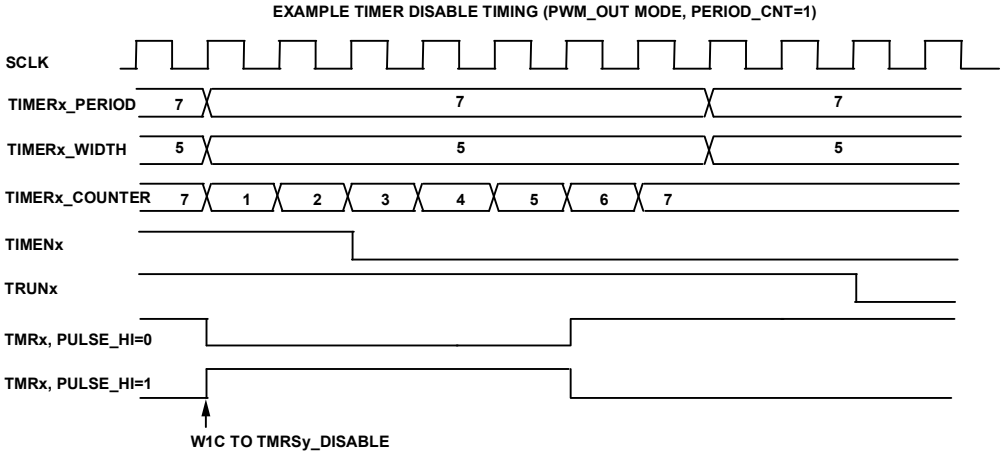

Figure 15-12. Timer Disable Timing

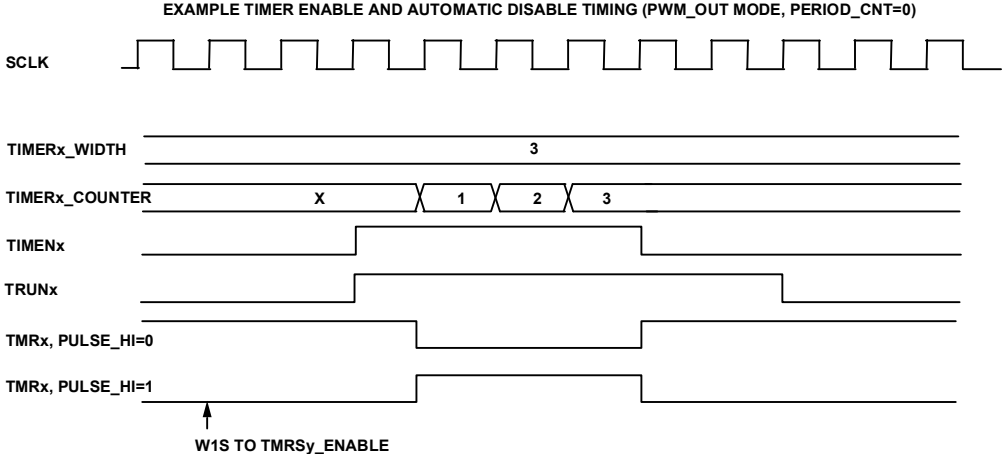

Figure 15-13. Timer Enable and Automatic Disable

Pulse Width Modulation (PWM_OUT) Mode

Setting the TMODE bits to 01 in the Timer Configuration (TIMERx_CONFIG) register enables PWM_OUT mode. In PWM_OUT mode, the timer TMRx pin is an output. The output can be disabled by setting the OUT_DIS bit in the Timer Configuration register. See Figure 15-14.

In PWM_OUT mode, the bits PULSE_HI, PERIOD_CNT, IRQ_ENA, OUT_DIS, CLK_SEL, EMU_RUN, and TOGGLE_HI enable orthogonal functionality. They may be set individually or in any combination, although some combinations are not useful (such as TOGGLE_HI = 1 with OUT_DIS = 1 or PERIOD_CNT = 0).

Once a timer has been enabled, the Timer Counter register is loaded with a starting value. If CLK_SEL=0, the Timer Counter starts at 0x1. If CLK_SEL=1, it is reset to 0x0s as in EXT_CLK mode. The timer counts upward to the value of the Timer Period register. For either setting of CLK_SEL, when the Timer Counter equals the Timer Period, the Timer Counter is reset to 0x1 on the next clock.

In PWM_OUT mode, the PERIOD_CNT bit controls whether the timer generates one pulse or many pulses. When PERIOD_CNT is cleared (PWM_OUT single pulse mode), the timer uses the TIMERx_WIDTH register, generates one asserting and one de-asserting edge, then generates an interrupt (if enabled) and stops. When PERIOD_CNT is set (PWM_OUT continuous pulse mode), the timer uses both the TIMERx_PERIOD and TIMERx_WIDTH registers and generates a repeating (and possibly modulated) waveform. It generates an interrupt (if enabled) at the end of each period and stops only after it is disabled. A setting of PERIOD_CNT = 0 counts to the end of the Width; a setting of PERIOD_CNT = 1 counts to the end of the Period.

Output Pad Disable

The output pin can be disabled in PWM_OUT mode by setting the OUT_DIS bit in the Timer Configuration register. The TMRx pin is then three-stated

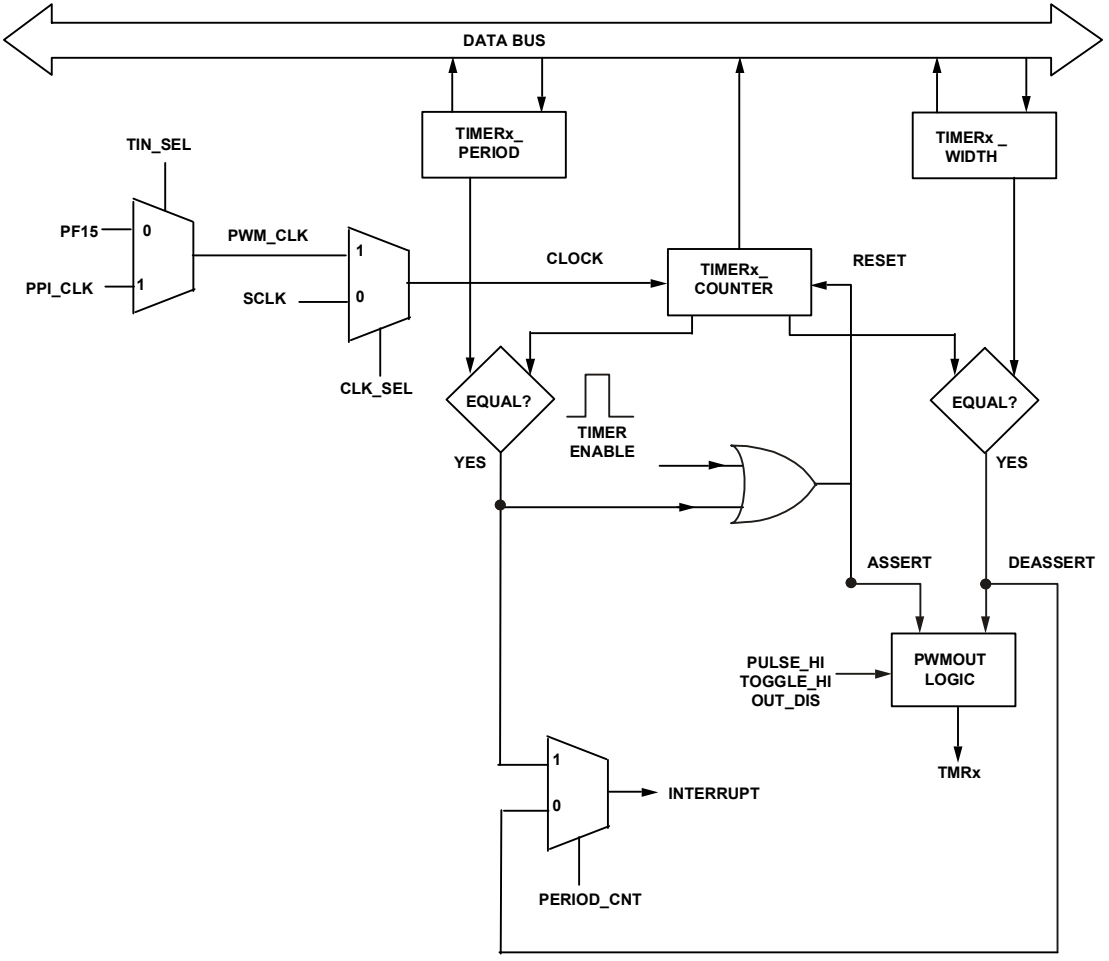

Figure 15-14. Timer Flow Diagram - PWM_OUT Mode

regardless of the setting of PULSE_HI and TOGGLE_HI. This can reduce power consumption when the output signal is not being used.

Single Pulse Generation

If the PERIOD_CNT bit is cleared, the PWM_OUT mode generates a single pulse on the TMRx pin. This mode can also be used to implement a precise delay. The pulse width is defined by the Timer Pulse Width register, and the Timer Period register is not used.

At the end of the pulse, the Timer Interrupt latch bit TIMILx gets set, and the timer is stopped automatically. If the PULSE_HI bit is set, an active high pulse is generated on the TMRx pin. If PULSE_HI is not set, the pulse is active low.

Pulse Width Modulation Waveform Generation

If the PERIOD_CNT bit is set, the internally clocked timer generates rectangular signals with well-defined period and duty cycle. This mode also generates periodic interrupts for real-time signal processing.

The 32-bit Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) registers are programmed with the values of the timer count period and pulse width modulated output.

When the timer is enabled in this mode, the TMRx pin is pulled to a deasserted state each time the Timer Counter equals the value of the Timer Pulse Width register, and the pin is asserted again when the period expires (or when the timer gets started).

To control the assertion sense of the TMRx pin, the PULSE_HI bit in the corresponding TIMERx_CONFIG register is used. For a low assertion level, clear this bit. For a high assertion level, set this bit. When the timer is disabled in PWM_OUT mode, the TMRx pin is driven to the deasserted level.

If enabled, a timer interrupt is generated at the end of each period. An interrupt service routine (ISR) must clear the Interrupt Latch bit (TIMILx) and might alter period and/or width values. In pulse width modulation (PWM) applications, the software needs to update period and pulse width values while the timer is running. When software updates either period or

pulse width registers, the new values are held by special buffer registers until the period expires. Then the new period and pulse width values become active simultaneously. New Timer Period and Timer Pulse Width register values are written while the old values are being used. The new values are loaded in to be used when the Timer Counter value equals the current Timer Period value. Reads from Timer Period and Timer Pulse Width registers return the old values until the period expires.

The TOVF_ERRx status bit signifies an error condition in PWM_OUT mode. The TOVF_ERRx bit is set if Period = 0 or Period = 1 at startup, or when the Timer Counter register rolls over. It is also set when the Timer Counter register rolls over if the Timer Pulse Width register is greater than or equal to the Timer Period register. The ERR_TYP bits are set when the TOVF_ERRx bit is set.

To generate the maximum frequency on the TMRx output pin, set the period value to 2 and the pulse width to 1. This makes TMRx toggle each SCLK clock, producing a duty cycle of 50%. The period may be programmed to any value from 2 to (2^{32} – 1), inclusive. The pulse width may be programmed to any value from 1 to (Period – 1), inclusive. When PERIOD_CNT = 0, the pulse width may be programmed to any value from 1 to (2^{32} – 1), inclusive.

Stopping the Timer in PWM_OUT Mode

In all PWM_OUT mode variants, the timer treats a disable operation (W1C to TMRSy_DISABLE) as a “stop is pending” condition. When disabled, it automatically completes the current waveform and then stops cleanly. This prevents truncation of the current pulse and unwanted PWM patterns at the TMRx pin. The processor can determine when the timer stops running by polling for the corresponding TRUNx bit in the TMRSy_STATUS register to read 0 or by waiting for the last interrupt (if enabled). Note the timer cannot be reconfigured (TIMERx_CONFIG cannot be written to a new value) until after the timer stops and TRUNx reads 0.

In PWM_OUT single pulse mode (PERIOD_CNT = 0), it is not necessary to write TMRSy_DISABLE to stop the timer. At the end of the pulse, the timer stops automatically, the corresponding bit in TMRSy_ENABLE (and TMRSy_DISABLE) is cleared, and the corresponding TRUNx bit is cleared. See Figure 15-13 on page 15-19. To generate multiple pulses, write a 1 to TMRSy_ENABLE, wait for the timer to stop, then write another 1 to TMRSy_ENABLE.

If necessary, the processor can force a timer in PWM_OUT mode to stop immediately. Do this by first writing a 1 to the corresponding bit in TMRSy_DISABLE, and then writing a 1 to the corresponding TRUNx bit in TMRSy_STATUS. This stops the timer whether the pending stop was waiting for the end of the current period (PERIOD_CNT = 1) or the end of the current pulse width (PERIOD_CNT = 0). This feature may be used to regain immediate control of a timer during an error recovery sequence.

Use this feature carefully, because it may corrupt the PWM pattern generated at the TMRx pin.

In PWM_OUT continuous pulse mode (PERIOD_CNT = 1), each timer samples its TIMENx bit at the end of each period. It stops cleanly at the end of the first period when TIMENx is low. This implies (barring any W1C to TRUNx) that a timer that is disabled and then re-enabled all before the end of the current period will continue to run as if nothing happened. Typically, software should disable a PWM_OUT timer and then wait for it to stop itself. The timer will always stop at the end of the first pulse when PERIOD_CNT = 0.

Externally clocked PWM_OUT

By default, the timer is clocked internally by SCLK. Alternatively, if the CLK_SEL bit in the Timer Configuration (TIMERx_CONFIG) register is set, then the timer is clocked by PWM_CLK. The PWM_CLK is by default input from the PF15 pin, but may be taken from the PPI_CLK pin when the timers are configured to work with the PPI. Different timers may receive different signals on their PWM_CLK inputs, depending on configuration. As

selected by the PERIOD_CNT bit, the PWM_OUT mode either generates pulse width modulation waveforms or generates a single pulse with pulse width defined by the TIMERx_WIDTH register.

When CLK_SEL is set, the counter resets to 0x0 at startup and increments on each rising edge of PWM_CLK. The TMRx pin transitions on rising edges of PWM_CLK. There is no way to select the falling edges of PWM_CLK. In this mode, the PULSE_HI bit controls only the polarity of the pulses produced. The timer interrupt may occur slightly before the corresponding edge on the TMRx pin (the interrupt occurs on an SCLK edge, the pin transitions on a later PWM_CLK edge). It is still safe to program new period and pulse width values as soon as the interrupt occurs. After a period expires, the counter rolls over to a value of 0x1.

The PWM_CLK clock waveform is not required to have a 50% duty cycle, but the minimum PWM_CLK clock low time is one SCLK period, and the minimum PWM_CLK clock high time is one SCLK period. This implies the maximum PWM_CLK clock frequency is SCLK/2.

The PF15 pin can only clock the timer when PF15 functions as an input pin. When any timer is in PWM_OUT mode with CLK_SEL = 1 and TIN_SEL = 0, then the PF15 bit in the FIO0_DIR register is ignored and PF15 is forced to be an input.

PULSE_HI Toggle Mode

The waveform produced in PWM_OUT mode with PERIOD_CNT = 1 normally has a fixed assertion time and a programmable de-assertion time (via the TIMERx_WIDTH register). When two timers are running synchronously by the same period settings, the pulses are aligned to the asserting edge as shown in Figure 15-15.

TOGGLE_HI mode enables control of the timing of both the asserting and de-asserting edges of the output waveform produced. The phase between the asserting edges of two timer outputs is programmable. The effective state of the PULSE_HI bit alternates every period. The adjacent active low

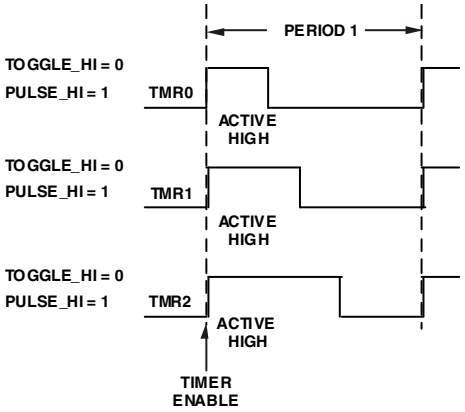

Figure 15-15. Timers With Pulses Aligned to Asserting Edge

and active high pulses, taken together, create two halves of a fully arbitrary rectangular waveform. The effective waveform is still active high when `PULSE_HI` is set and active low when `PULSE_HI` is cleared. `TOGGLE_HI` has no effect unless the mode is `PWM_OUT` and `PERIOD_CNT = 1`.

In `TOGGLE_HI` mode when `PULSE_HI` is set, an active *low* pulse is generated in the first, third, and all odd-numbered periods, and an active *high* pulse is generated in the second, fourth, and all even-numbered periods. When `PULSE_HI` is cleared, an active *high* pulse is generated in the first, third, and all odd-numbered periods, and an active *low* pulse is generated in the second, fourth, and all even-numbered periods.

The deasserted state at the end of one period matches the asserted state at the beginning of the next period, so the output waveform only transitions when `Count = Pulse Width`. The net result is an output waveform pulse that repeats every two counter periods and is centered around the end of the first period (or the start of the second period).

Figure 15-16 shows an example with three timers running with the same period settings. When software does not alter the PWM settings at

run-time, the duty cycle is 50%. The values of the TIMERx_WIDTH registers control the phase between the signals.

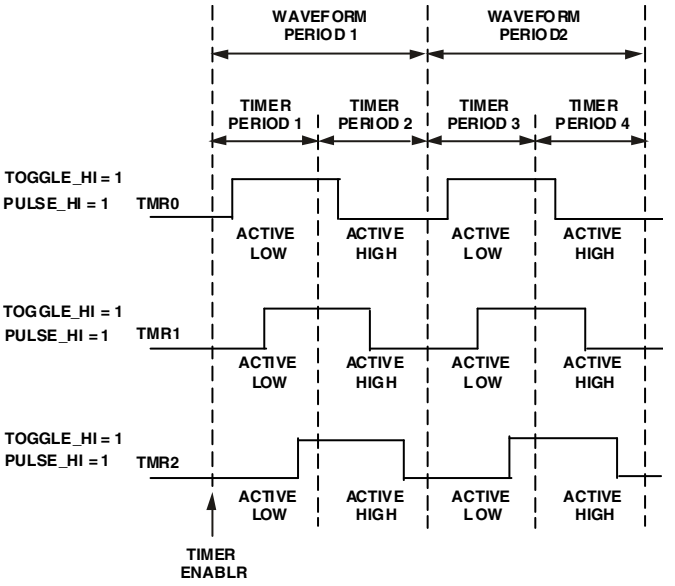

Figure 15-16. Three Timers With Same Period Settings

Similarly, two timers can generate non-overlapping clocks. Center align the pulses but invert the signal polarity for one of the timers, as shown in Figure 15-17.

When TOGGLE_HI = 0, software updates the Timer Period and Timer Pulse Width registers once per waveform period. When TOGGLE_HI = 1, software updates the Timer Period and Timer Pulse Width registers twice per waveform period with values that are half as large. In odd-numbered periods, write (Period - Width), instead of Width, to the Timer Pulse Width register in order to obtain center-aligned pulses.

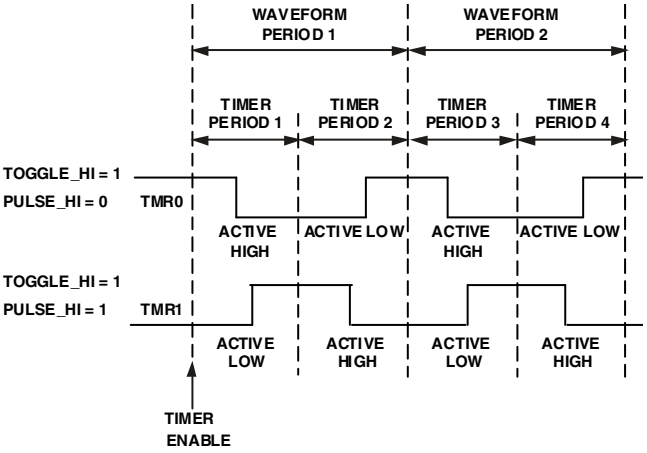

Figure 15-17. Two Timers With Non-Overlapping Clocks

For example, if the pseudocode when `TOGGLE_HI = 0` is:

```
int period, width ;
for (;;) {
     period = generate_period(...) ;
     width = generate_width(...) ;

     waitfor (interrupt) ;

     write(TIMERx_PERIOD, period) ;
     write(TIMERx_WIDTH, width) ;
}
```

Then when `TOGGLE_HI = 1`, the pseudocode would be:

```
int period, width ;
int per1, per2, wid1, wid2 ;

for (;;) {
     period = generate_period(...) ;
```

```
        width = generate_width(...) ;

        per1 = period/2 ;
        wid1 = width/2 ;

        per2 = period/2 ;
        wid2 = width/2 ;

        waitfor (interrupt) ;

        write(TIMERx_PERIOD, per1) ;
        write(TIMERx_WIDTH, per1 - wid1) ;

        waitfor (interrupt) ;

        write(TIMERx_PERIOD, per2) ;
        write(TIMERx_WIDTH, wid2) ;

}
```

Note that the pulses produced do not need to be symmetric as shown in this example (`wid1` does not need to equal `wid2`). The period can be offset to adjust the phase of the pulses produced (`per1` does not need to equal `per2`).

The slave timer enable bit (`TRUN` bit in the `TMRSy_STATUS` register) is updated only at the end of even-numbered periods in `TOGGLE_HI` mode. When `TMRSy_DISABLE` is written to 1, the current pair of counter periods (one waveform period) completes before the timer is disabled.

As when `TOGGLE_HI = 0`, errors are reported if `TIMERx_WIDTH >= TIMERx_PERIOD`, `TIMERx_PERIOD = 0`, or `TIMERx_PERIOD = 1`.

Pulse Width Count and Capture (WDTH_CAP) Mode

In WDTH_CAP mode, the TMRx pin is an input pin. The internally clocked timer is used to determine the period and pulse width of externally applied rectangular waveforms. Setting the TMODE field to b#10 in the TIMERx_CONFIG (Timer Configuration register) enables this mode. See Figure 15-18.

When enabled in this mode, the timer resets the count in the TIMERx_COUNTER register to 0x0000 0001 and does not start counting until it detects a leading edge on the TMRx pin.

When the timer detects the first leading edge, it starts incrementing. When it detects a trailing edge of a waveform, the timer captures the current 32-bit value of the TIMERx_COUNTER register into the width buffer register. At the next leading edge, the timer transfers the current 32-bit value of the TIMERx_COUNTER register into the period buffer register. The count register is reset to 0x0000 0001 again, and the timer continues counting and capturing until it is disabled.

In this mode, software can measure both the pulse width and the pulse period of a waveform. To control the definition of “leading edge” and “trailing edge” of the TMRx pin, the PULSE_HI bit in the TIMERx_CONFIG register is set or cleared. If the PULSE_HI bit is cleared, the measurement is initiated by a falling edge, the Timer Counter register is captured to the Timer Pulse Width buffer register on the rising edge, and the Timer Period is captured on the next falling edge. When the PULSE_HI bit is set, the measurement is initiated by a rising edge, the Timer Counter register is captured to the Timer Pulse Width buffer register on the falling edge, and the Timer Period is captured on the next rising edge.

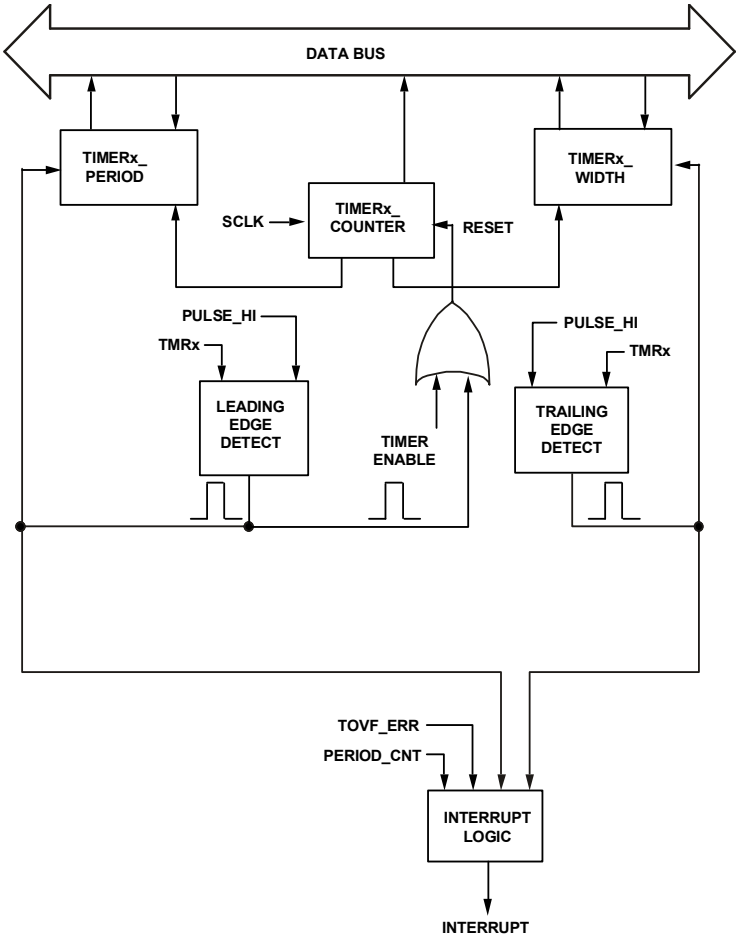

Figure 15-18. Timer Flow Diagram - WDTH_CAP Mode

In WDTH_CAP mode, these three events always occur at the same time as one unit:

1. The TIMERx_PERIOD register is updated from the period buffer register.
2. The TIMERx_WIDTH register is updated from the width buffer register.
3. The Timer Interrupt latch bit (TIMILx) gets set (if enabled) but does not generate an error.

The PERIOD_CNT bit in the TIMERx_CONFIG register controls the point in time at which this set of transactions is executed. Taken together, these three events are called a measurement report. The Timer Counter Overflow error latch bit (TOVF_ERRx) does not get set at a measurement report. A measurement report occurs at most once per input signal period.

The current timer counter value is always copied to the width buffer and period buffer registers at the trailing and leading edges of the input signal, respectively, but these values are not visible to software. A measurement report event samples the captured values into visible registers and sets the timer interrupt to signal that TIMERx_PERIOD and TIMERx_WIDTH are ready to be read. When the PERIOD_CNT bit is set, the measurement report occurs just after the period buffer register captures its value (at a leading edge). When the PERIOD_CNT bit is cleared, the measurement report occurs just after the width buffer register captures its value (at a trailing edge).

If the PERIOD_CNT bit is set and a leading edge occurred (Figure 15-19), then the TIMERx_PERIOD and TIMERx_WIDTH registers report the pulse period and pulse width measured in the period that just ended. If the PERIOD_CNT bit is cleared and a trailing edge occurred (Figure 15-20), then the TIMERx_WIDTH register reports the pulse width measured in the pulse that just ended, but the TIMERx_PERIOD register reports the pulse period measured at the end of the previous period.

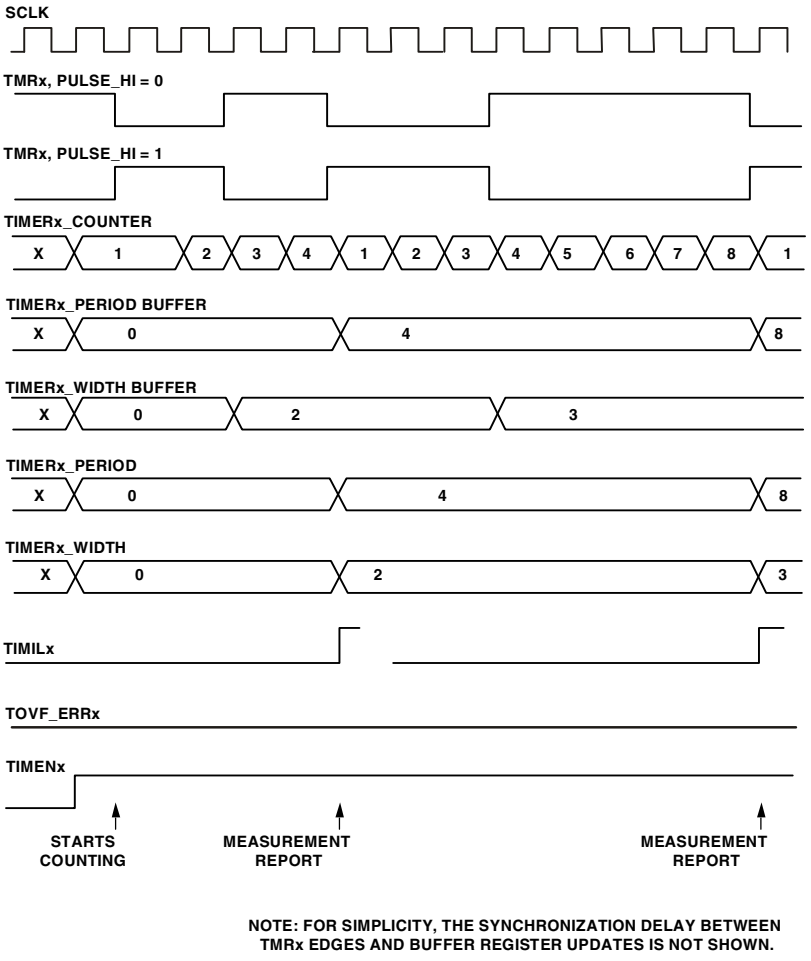

Figure 15-19. Example of Period Capture Measurement Report Timing (WDTH_CAP mode, PERIOD_CNT = 1)

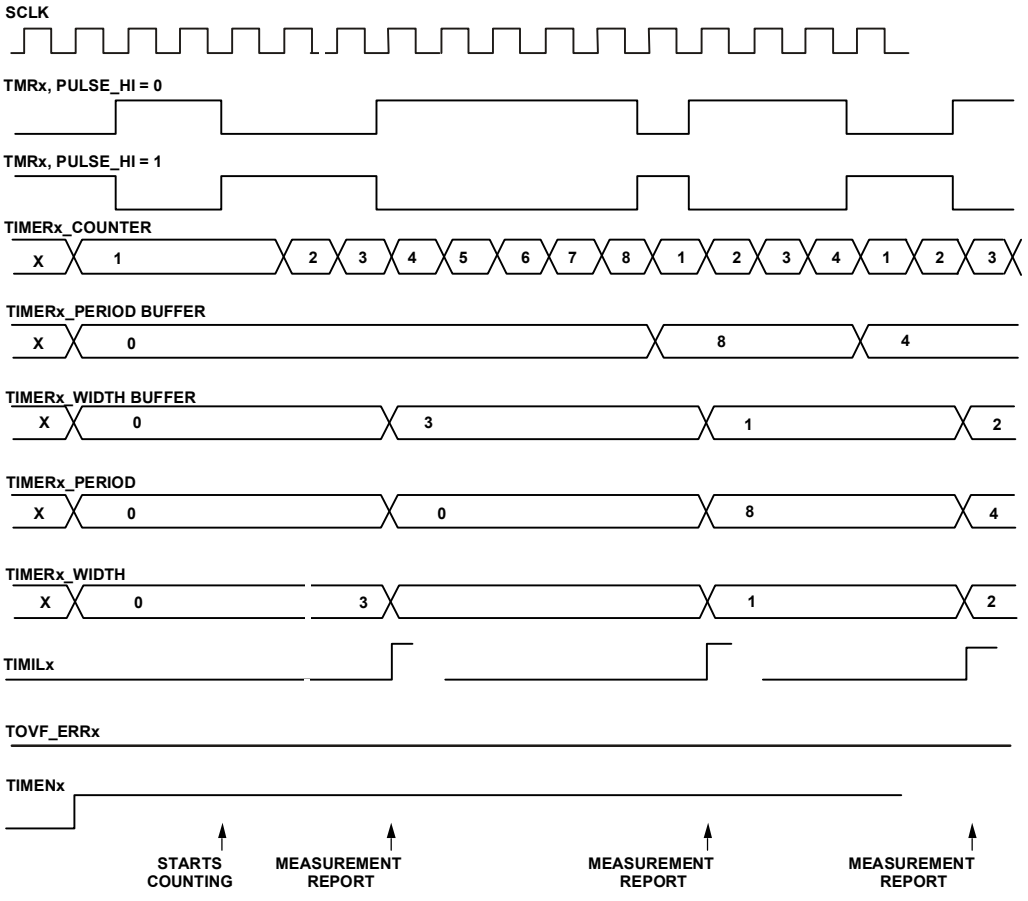

Figure 15-20. Example of Width Capture Measurement Report Timing (WDTH_CAP mode, PERIOD_CNT = 0)

If the PERIOD_CNT bit is cleared and the first trailing edge occurred, then the first period value has not yet been measured at the first measurement report, so the period value is not valid. Reading the TIMERx_PERIOD value in this case returns 0, as shown in Figure 15-20. To measure the pulse

width of a waveform that has only one leading edge and one trailing edge, set PERIOD_CNT = 0. If PERIOD_CNT = 1 for this case, no period value is captured in the period buffer register. Instead, an error report interrupt is generated (if enabled) when the counter range is exceeded and the counter wraps around. In this case, both TIMERx_WIDTH and TIMERx_PERIOD read 0 (because no measurement report occurred to copy the value captured in the width buffer register to TIMERx_WIDTH). See the first interrupt in Figure 15-21.

When using the PERIOD_CNT = 0 mode described above to measure the width of a single pulse, it is recommended to disable the timer after taking the interrupt that ends the measurement interval. If desired, the timer can then be re-enabled as appropriate in preparation for another measurement. This procedure prevents the timer from free-running after the width measurement and logging errors generated by the timer count overflowing.

A timer interrupt (if enabled) is generated if the Timer Counter register wraps around from 0xFFFF FFFF to 0 in the absence of a leading edge. At that point, the TOVF_ERRx bit in the TMRSy_STATUS register and the ERR_TYP bits in the TIMERx_CONFIG register are set, indicating a count overflow due to a period greater than the counter's range. This is called an error report. When a timer generates an interrupt in WDTH_CAP mode, either an error has occurred (an error report) or a new measurement is ready to be read (a measurement report), but never both at the same time. The TIMERx_PERIOD and TIMERx_WIDTH registers are never updated at the time an error is signaled. Refer to Figure 15-21 and Figure 15-22 for more information.

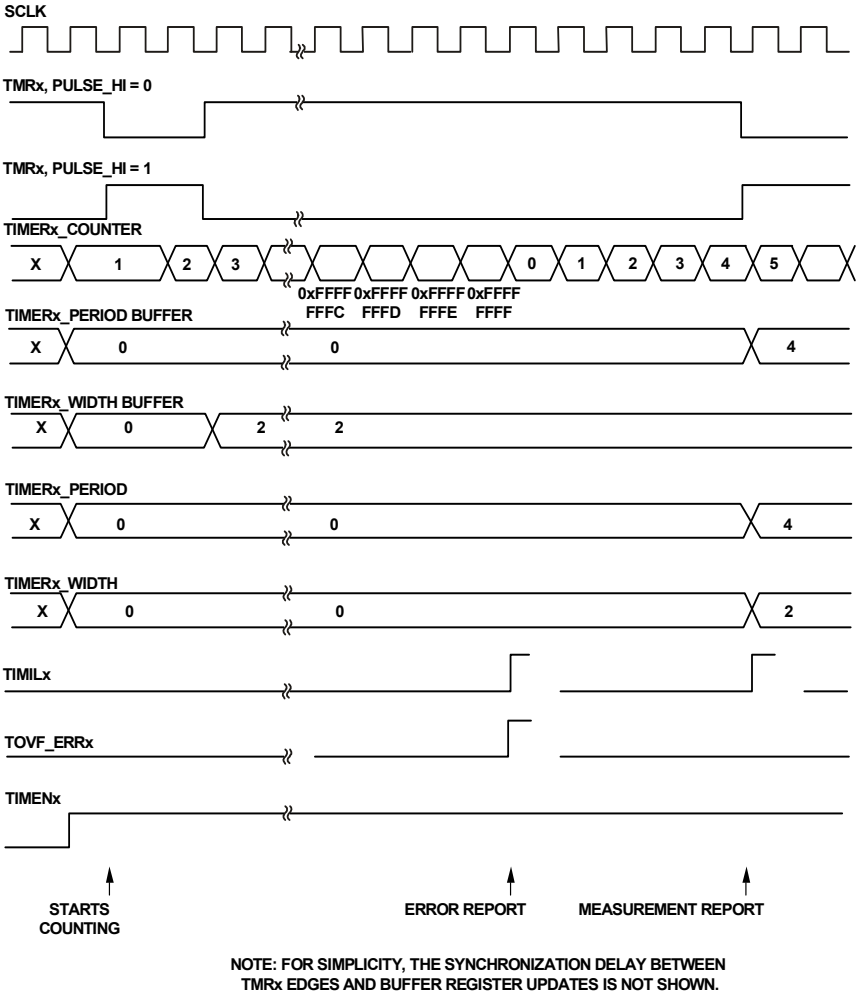

Figure 15-21. Example Timing for Period Overflow Followed by Period Capture (WDTH_CAP mode, PERIOD_CNT = 1)

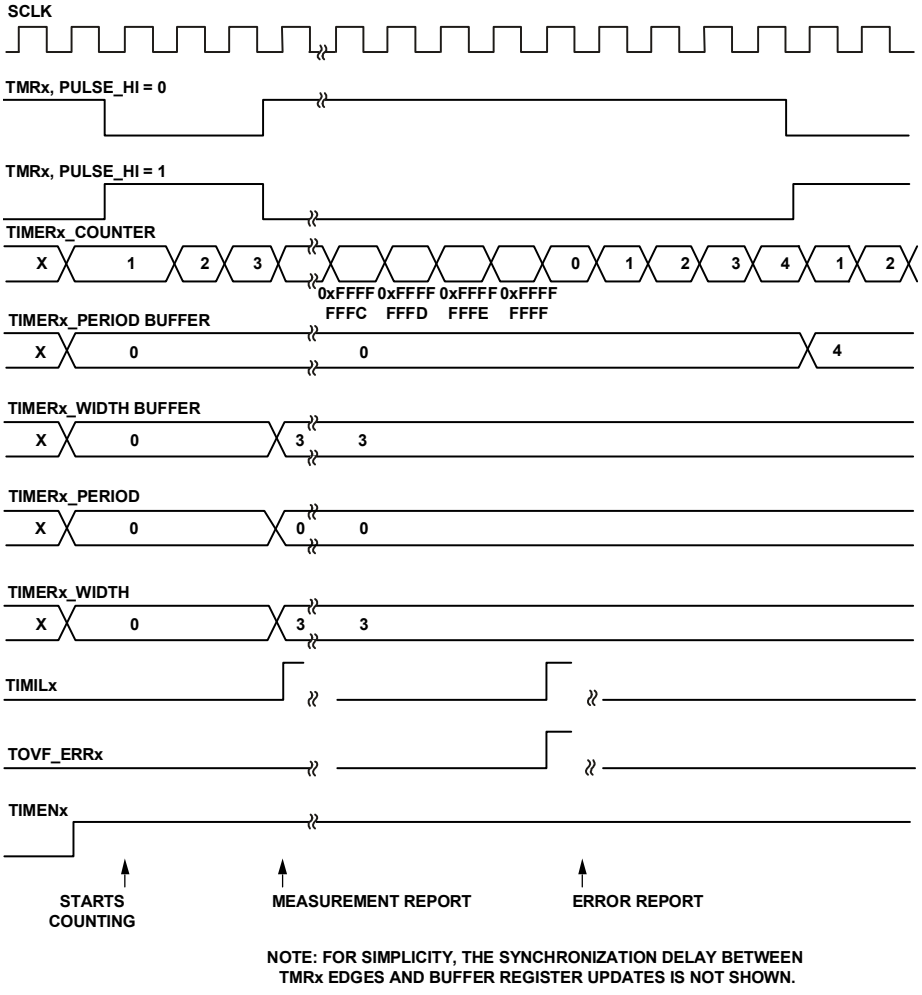

Figure 15-22. Example Timing for Width Capture Followed by Period Overflow (WDTH_CAP mode, PERIOD_CNT = 0)

Both TIMILx and TOVF_ERRx are sticky bits, and software has to explicitly clear them. If the timer overflowed and PERIOD_CNT = 1, neither the TIMERx_PERIOD nor the TIMERx_WIDTH register were updated. If the timer overflowed and PERIOD_CNT = 0, the TIMERx_PERIOD and TIMERx_WIDTH

registers were updated only if a trailing edge was detected at a previous measurement report.

Software can count the number of error report interrupts between measurement report interrupts to measure input signal periods longer than 0xFFFF FFFF. Each error report interrupt adds a full 2^{32} SCLK counts to the total for the period, but the width is ambiguous. For example, in Figure 15-21 the period is 0x1 0000 0004 but the pulse width could be either 0x0 0000 0002 or 0x1 0000 0002.

The waveform applied to the TMRx pin is not required to have a 50% duty cycle, but the minimum TMRx low time is one SCLK period and the minimum TMRx high time is one SCLK period. This implies the maximum TMRx input frequency is SCLK/2 with a 50% duty cycle. Under these conditions, the WDTH_CAP mode timer would measure Period = 2 and Pulse Width = 1.

Autobaud Mode

Any one of the three timers may provide autobaud detection for the Universal Asynchronous Receiver/Transmitter (UART). The Timer Input Select (TIN_SEL) bit in the TIMERx_CONFIG register causes the timer to sample the UART port receive data (RX) pin instead of the TMRx pin when enabled for WDTH_CAP mode.

Do not enable the UART until after autobaud detection is complete.

A software routine can detect the pulse widths of serial stream bit cells. Because the sample base of the timers is synchronous with the UART operation—all derived from the Phase Locked Loop (PLL) clock—the pulse widths can be used to calculate the baud rate divider for the UART.

DIVISOR = ((TIMERx_WIDTH) / (16 × Number of captured UART bits))

In order to increase the number of timer counts and therefore the resolution of the captured signal, it is recommended not to measure just the pulse width of a single bit, but to enlarge the pulse of interest over more bits. Typically a NULL character (ASCII 0x00) is used in autobaud detection, as shown in Figure 15-23.

Figure 15-23. Autobaud Detection Character 0x00

Because the example frame in Figure 15-23 encloses 8 data bits and 1 start bit, apply the formula:

```
DIVISOR = TIMERx_WIDTH/(16 × 9)
```

Real UART RX signals often have asymmetrical falling and rising edges, and the sampling logic level is not exactly in the middle of the signal voltage range. At higher bit rates, such pulse width-based autobaud detection might not return adequate results without additional analog signal conditioning. Measuring signal periods works around this issue and is strongly recommended.

For example, predefine ASCII character “@” (40h) as an autobaud detection byte and measure the period between two subsequent falling edges. As shown in Figure 15-24, measure the period between the falling edge of the start bit and the falling edge after bit 6. Since this period encloses 8 bits, apply the formula:

```
DIVISOR = TIMERx_PERIOD/(16 × 8)
```

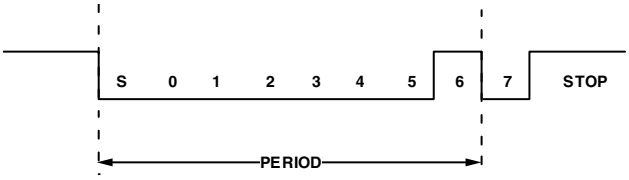

Figure 15-24. Autobaud Detection Character 0x40

External Event (EXT_CLK) Mode

In EXT_CLK mode, the TMRx pin is an input. The timer works as a counter clocked by an external source, which can also be asynchronous to the system clock. The current count in TIMERx_COUNTER represents the number of leading edge events detected. Setting the TMODE field to b#11 in the TIMERx_CONFIG register enables this mode. The TIMERx_PERIOD register is programmed with the value of the maximum timer external count.

The waveform applied to the TMRx pin is not required to have a 50% duty cycle, but the minimum TMRx low time is one SCLK period, and the minimum TMRx high time is one SCLK period. This implies the maximum TMRx input frequency is SCLK/2.

Period may be programmed to any value from 1 to $(2^{32} – 1)$, inclusive.

After the timer has been enabled, it resets the Timer Counter register to 0x0 and then waits for the first leading edge on the TMRx pin. This edge causes the Timer Counter register to be incremented to the value 0x1. Every subsequent leading edge increments the count register. After reaching the period value, the TIMILx bit is set, and an interrupt is generated. The next leading edge reloads the Timer Counter register again with 0x1. The timer continues counting until it is disabled. The PULSE_HI bit determines whether the leading edge is rising (PULSE_HI set) or falling (PULSE_HI cleared).

The configuration bits, TIN_SEL and PERIOD_CNT, have no effect in this mode. The TOVF_ERRx and ERR_TYP bits are set if the Timer Counter

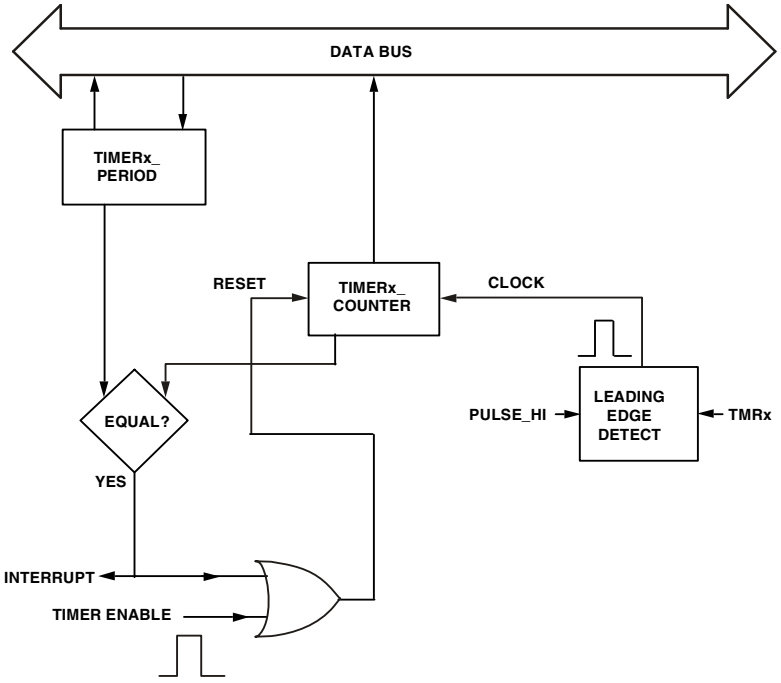

Figure 15-25. Timer Flow Diagram, EXT_CLK Mode

register wraps around from 0xFFFF FFFF to 0 or if `Period = 0` at startup or when the Timer Counter register rolls over (from `Count = Period` to `Count = 0x1`). The Timer Pulse Width register is unused.

Using the Timers With the PPI

Up to two timers are used to generate frame sync signals for certain PPI modes. For detailed instructions on how to configure the timers for use with the PPI, refer to "Frame Synchronization in GP Modes" on page 11-29 of the PPI chapter.

Interrupts

Each of the twelve timers can generate a single interrupt. The twelve resulting interrupt signals are routed to the System Interrupt Control block for prioritization and masking (Figure 15-26). The Timer Status (TMRSy_STATUS) registers latch the timer interrupts to provide a means for software to determine the interrupt source. These bits are W1C and must be cleared prior to a RETI to assure that the interrupt is not reissued.

To enable interrupt generation, set the IRQ_ENA bit and unmask the interrupt source in the System Interrupt Mask (SICx_IMASK) register. To poll the TIMILx bit without interrupt generation, set IRQ_ENA but leave the interrupt masked. If enabled by IRQ_ENA, interrupt requests are also generated by error conditions.

The system interrupt controller enables flexible interrupt handling. All timers may or may not share the same interrupt channel, so that a single interrupt routine services more than one timer. In PWM mode, more timers may run with the same period settings and issue their interrupt requests simultaneously.

If interrupts are enabled, make sure that the interrupt service routine (ISR) clears the TIMILx bit in the TMRSy_STATUS register before the RTI instruction executes. This ensures that the interrupt is not reissued. Remember that writes to system registers are delayed. If only a few instructions separate the TIMILx clear command from the RTI instruction, an extra SSYNC instruction may be inserted. In EXT_CLK mode, reset the TIMILx bit in the TMRSy_STATUS register at the very beginning of the ISR to avoid missing any timer events.

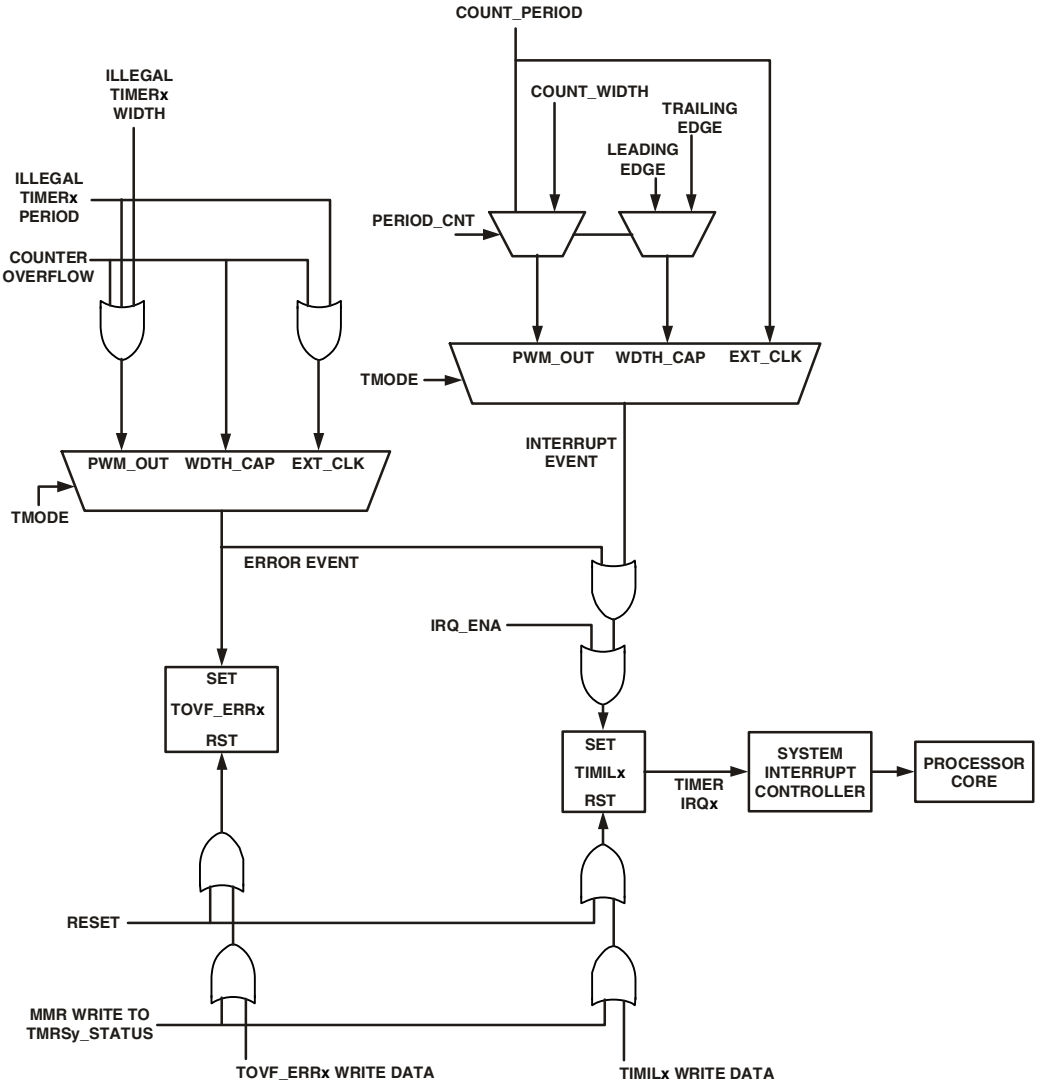

Figure 15-26. Timers Interrupt Structure

Illegal States

These definitions are used in Table 15-1.

- Startup. The first clock period during which the timer counter is running after the timer is enabled by writing TMRSy_ENABLE.
- Rollover. The time when the current count matches the value in TIMERx_PERIOD and the counter is reloaded with the value 1.
- Overflow. The timer counter was incremented instead of doing a rollover when it was holding the maximum possible count value of 0xFFFF FFFF. The counter does not have a large enough range to express the next greater value and so erroneously loads a new value of 0x0000 0000.
- Unchanged. No new error.

  When ERR_TYP is unchanged, it displays the previously reported error code or 00 if there has been no error since this timer was enabled.

  When TOVF_ERR is unchanged, it reads 0 if there has been no error since this timer was enabled, or if software has performed a W1C to clear any previous error. If a previous error has not been acknowledged by software, TOVF_ERR reads 1.

Software should read TOVF_ERR to check for an error. If TOVF_ERR is set, software can then read ERR_TYP for more information. Once detected, software should W1C TOVF_ERR to acknowledge the error.

The following table can be read as: “In mode __ at event __, if TIMERx_PERIOD is __ and TIMERx_WIDTH is __, then ERR_TYP is __ and TOVF_ERR is __.”

Startup error conditions do not prevent the timer from starting. Similarly, overflow and rollover error conditions do not stop the timer. Illegal cases may cause unwanted behavior of the TMRx pin.

Table 15-1. Overview of Illegal States

Mode	Event	TIMERx_ PERIOD	TIMERx_ WIDTH	ERR_TYP	TOVF_ERR
PWM_OUT, PERIOD_ CNT = 1	Startup (No boundary condition tests performed on TIMERx_ WIDTH)	== 0	Anything	b#10	Set
		== 1	Anything	b#10	Set
		>= 2	Anything	Unchanged	Unchanged
	Rollover	== 0	Anything	b#10	Set
		== 1	Anything	b#11	Set
		>= 2	== 0	b#11	Set
		>= 2	< TIMERx_ PERIOD	Unchanged	Unchanged
		>= 2	>= TIMERx_ PERIOD	b#11	Set
	Overflow, not possible unless there is also another error, such as TIMERx_ PERIOD == 0.	Anything	Anything	b#01	Set

Table 15-1. Overview of Illegal States (Cont'd)

Mode	Event	TIMERx_ PERIOD	TIMERx_ WIDTH	ERR_TYP	TOVF_ERR
PWM_OUT, PERIOD_ CNT = 0	Startup	Anything	== 0	b#01	Set
		This case is not detected at startup, but results in an overflow error once the counter counts through its entire range.			
		Anything	>= 1	Unchanged	Unchanged
	Rollover	Rollover is not possible in this mode.			
	Overflow, not possible unless there is also another error, such as TIMERx_ WIDTH == 0.	Anything	Anything	b#01	Set
WDTH_CAP	Startup	TIMERx_PERIOD and TIMERx_WIDTH are read-only in this mode, no error possible.			
	Rollover	TIMERx_PERIOD and TIMERx_WIDTH are read-only in this mode, no error possible.			
	Overflow	Anything	Anything	b#01	Set
EXT_CLK	Startup	== 0	Anything	b#10	Set
		>= 1	Anything	Unchanged	Unchanged
	Rollover	== 0	Anything	b#10	Set
		>= 1	Anything	Unchanged	Unchanged
	Overflow, not possible unless there is also another error, such as TIMERx_ PERIOD == 0.	Anything	Anything	b#01	Set

Summary

Table 15-2 summarizes the control bit and register usage in each timer mode.

Table 15-2. Control Bit and Register Usage Chart

Bit / Register	PWM_OUT Mode	WDTH_CAP Mode	EXT_CLK Mode
TMRSy_ENABLE	1 - Enable timer. 0 - No effect.	1 - Enable timer. 0 - No effect.	1 - Enable timer. 0 - No effect.
TMRSy_DISABLE	1 - Disable timer at end of period. 0 - No effect.	1 - Disable timer. 0 - No effect.	1 - Disable timer. 0 - No effect.
TMODE	b#01	b#10	b#11
PULSE_HI	1 - Generate high Width. 0 - Generate low Width.	1 - Measure high Width. 0 - Measure low Width.	1 - Count rising edges. 0 - Count falling edges.
PERIOD_CNT	1 - Generate PWM. 0 - Single Width pulse.	1 - Interrupt after measuring Period. 0 - Interrupt after measuring Width.	Unused.
IRQ_ENA	1 - Enable interrupt. 0 - Disable interrupt.	1 - Enable interrupt. 0 - Disable interrupt.	1 - Enable interrupt. 0 - Disable interrupt.
TIN_SEL	Depends on CLK_SEL: If CLK_SEL = 1, 1 - Count PPI_CLKs 0 - Count PF15 clocks If CLK_SEL = 0, Unused	1 - Select RX input. 0 - Select TMRx input.	Unused.
OUT_DIS	1 - Disable TMRx pin. 0 - Enable TMRx pin.	Unused.	Unused.
CLK_SEL	1 - PWM_CLK clocks timer 0 - SCLK clocks timer	Unused.	Unused.

Table 15-2. Control Bit and Register Usage Chart (Cont'd)

Bit / Register	PWM_OUT Mode	WDTH_CAP Mode	EXT_CLK Mode
TOGGLE_HI	1 - One waveform period every two counter periods. 0 - One waveform period every one counter period.	Unused.	Unused.
ERR_TYP	Reports b#00, b#01, b#10, or b#11, as appropriate.	Reports b#00 or b#01, as appropriate.	Reports b#00, b#01, or b#10, as appropriate.
EMU_RUN	0 - Halt during emulation 1 - Count during emulation	0 - Halt during emulation 1 - Count during emulation	0 - Halt during emulation 1 - Count during emulation
TMR Pin	Depends on OUT_DIS: 1 - Three-state. 0 - Output.	Depends on TIN_SEL: 1 - Unused. 0 - Input.	Input.
Period	R/W: Period value.	RO: Period value.	R/W: Period value.
Width	R/W: Width value.	RO: Width value.	Unused.
Counter	RO: Counts up on SCLK or PWM_CLK	RO: Counts up on SCLK.	RO: Counts up on event.
TRUN	Read: Timer slave enable status. Write: 1 - Stop timer if disabled. 0 - No effect.	Read: Timer slave enable status. Write: 1 - No effect. 0 - No effect.	Read: Timer slave enable status. Write: 1 - No effect. 0 - No effect.

Table 15-2. Control Bit and Register Usage Chart (Cont'd)

Bit / Register	PWM_OUT Mode	WDTH_CAP Mode	EXT_CLK Mode
TOVF_ERR	Set at startup or rollover if Period = 0 or 1. Set at rollover if Width >= Period. Set if Counter wraps.	Set if Counter wraps.	Set if Counter wraps or set at startup or rollover if Period = 0.
IRQ	Depends on IRQ_ENA: 1 - Set when TOVF_ERR set or when Counter equals Period and PERIOD_CNT = 1 or when Counter equals Width and PERIOD_CNT = 0. 0 - Not set.	Depends on IRQ_ENA: 1 - Set when TOVF_ERR set or when Counter captures Period and PERIOD_CNT = 1 or when Counter captures Width and PERIOD_CNT = 0. 0 - Not set.	Depends on IRQ_ENA: 1 - Set when Counter equals Period or TOVF_ERR set. 0 - Not set.

Core Timer

The core timer (one for each core) is a programmable interval timer which can generate periodic interrupts. The core timer runs at the core clock (CCLK) rate. The timer includes four core MMRs, the Timer Control register (TCNTL), the Timer Count register (TCOUNT), the Timer Period register (TPERIOD), and the Timer Scale register (TSCALE).

Figure 15-27 provides a block diagram of the core timer.

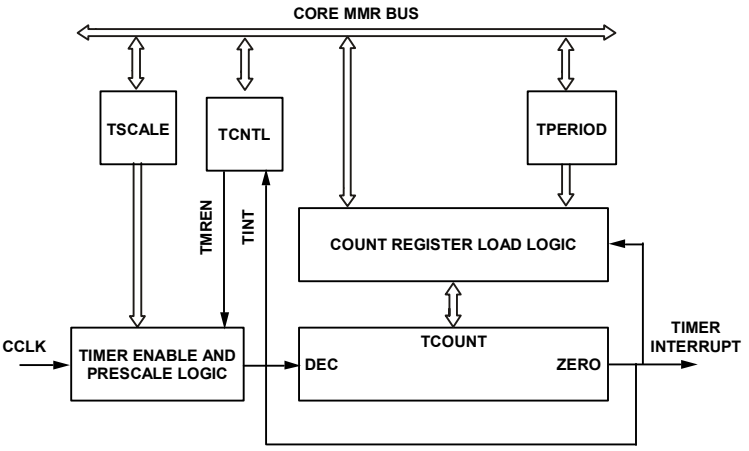

Figure 15-27. Core Timer Block Diagram

Core Timer Control Register (TCNTL)

When the timer is enabled by setting the TMREN bit in the TCNTL register (Figure 15-28), the TCOUNT register is decremented once every TSCALE+1 number of clock cycles. When the value of the TCOUNT register reaches 0, an interrupt is generated and the TINT bit is set in the TCNTL register. If the TAUTORLD bit in the TCNTL register is set, then the TCOUNT register is reloaded with the contents of the TPERIOD register and the count begins again.

The TINT bit in the TCNTL register indicates that an interrupt has been generated. Note that this is not a W1C bit. Write a 0 to clear it. However, the write is optional. It is not required to clear interrupt requests. The core time module does not provide any further interrupt enable bit. When the timer is enabled, interrupts can be masked in the CEC controller.

The core timer can be put into low power mode by clearing the TMPWR bit in the TCNTL register. Before using the timer, set the TMPWR bit. This restores clocks to the timer unit. When TMPWR is set, the core timer may then be enabled by setting the TMREN bit in the TCNTL register.

Hardware behavior is undefined if TMREN is set when TMPWR = 0.

Core Timer Control Register (TCNTL)

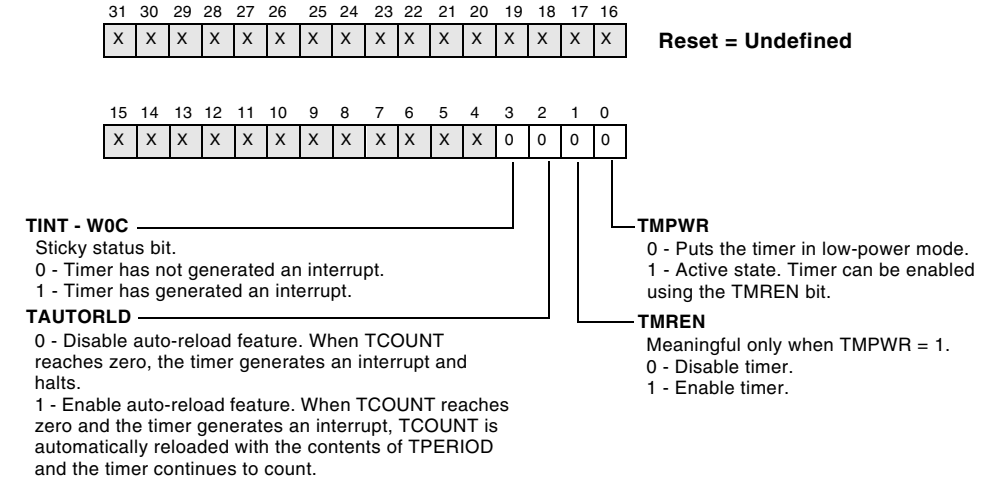

Figure 15-28. Core Timer Control Register

Core Timer Count Register (TCOUNT)

This register (Figure 15-29) decrements once every TSCALE + 1 clock cycles. When the value of TCOUNT reaches 0, an interrupt is generated and the TINT bit of the TCNTL register is set.

Core Timer Count Register (TCOUNT)

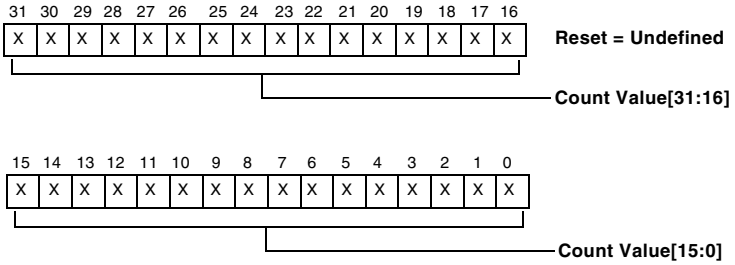

Figure 15-29. Core Timer Count Register

Core Timer Period Register (TPERIOD)

To ensure that there is valid data in the TPERIOD register (Figure 15-30), the TPERIOD and TCOUNT registers are initialized simultaneously on the first write to either register. If a different value is desired for the first count period, write the data to TCOUNT after writing to TPERIOD.

When auto-reload is enabled, the TCOUNT register is reloaded with the value of TPERIOD whenever TCOUNT reaches 0.

Core Timer Period Register (TPERIOD)

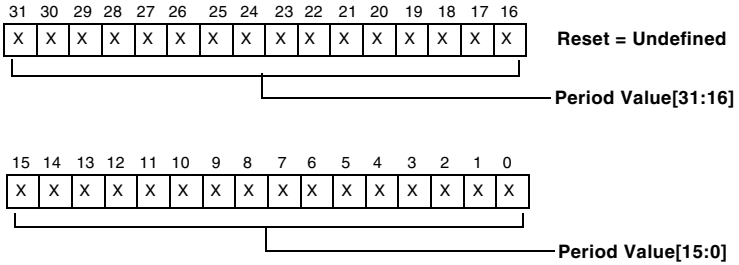

Figure 15-30. Core Timer Period Register

Core Timer Scale Register (TSCALE)

The TSCALE register (Figure 15-31) stores the scaling value that is one less than the number of cycles between decrements of TCOUNT. For example, if the value in the TSCALE register is zero, the counter register decrements once every clock cycle. If TSCALE is 1, the counter decrements once every 2 cycles.

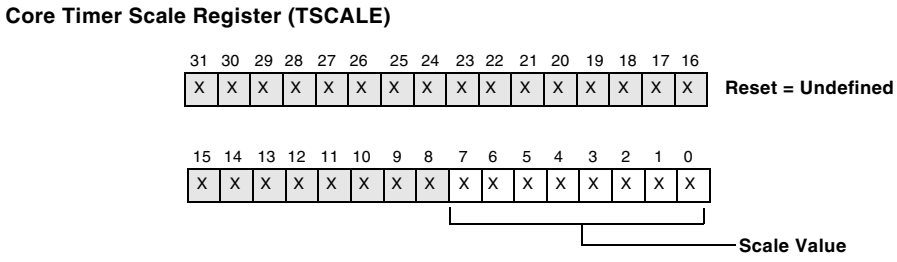

Figure 15-31. Core Timer Scale Register

Watchdog Timer

The processor includes a 32-bit timer for each core that can be used to implement a software watchdog function. A software watchdog can improve system reliability by generating an event to the processor core if the timer expires before being updated by software. Depending on how the watchdog timer is programmed, the event that is generated may be a reset, a nonmaskable interrupt, or a general-purpose interrupt. The watchdog timer is clocked by the system clock (SCLK).

Each core has its own watchdog timer register set, for example, WDOGA_CNT for core A, WDOGB_CNT for core B. They have identical definitions. In the following description, WDOG_CNT is used regardless whether it is for core A or core B. Other registers are represented in the same way.

Watchdog Timer Operation

To use the watchdog timer:

1. Set the count value for the watchdog timer by writing the count value into the Watchdog Count register (WDOG_CNT). Note that loading the WDOG_CNT register while the Watchdog timer is not enabled will also pre-load the WDOG_STAT register.
2. In the Watchdog Control register (WDOG_CTL), select the event to be generated upon timeout.
3. Enable the watchdog timer in WDOG_CTL. The watchdog timer then begins counting down, decrementing the value in the WDOG_STAT register. When the WDOG_STAT reaches 0, the programmed event is generated. To prevent the event from being generated, software must reload the count value from WDOG_CNT to WDOG_STAT by executing a write (of any value) to WDOG_STAT, or must disable the watchdog timer in WDG_CTL before the watchdog timer expires.

Watchdog Count Register (WDOG_CNT)

The Watchdog Count register (WDOG_CNT) holds the 32-bit unsigned count value. The WDOG_CNT register (Figure 15-32) must be accessed with 32-bit read/writes only.

The Watchdog Count register holds the programmable count value. A valid write to the Watchdog Count register also pre-loads the Watchdog counter. For added safety, the Watchdog Count register can be updated only when the Watchdog timer is disabled. A write to the Watchdog Count register while the timer is enabled does not modify the contents of this register.

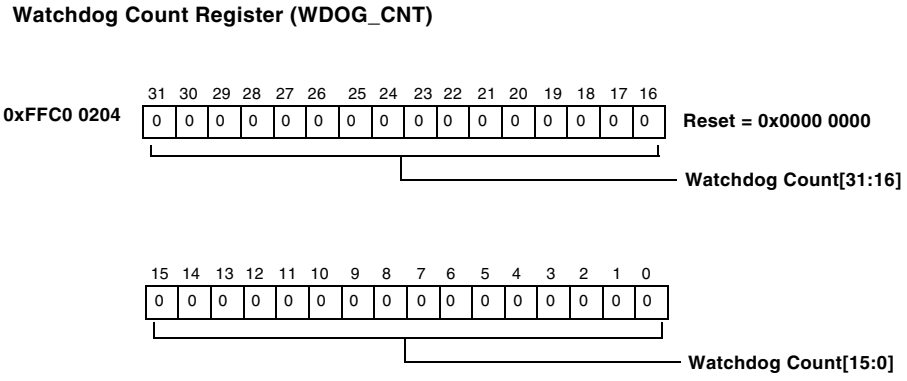

Figure 15-32. Watchdog Count Register

Watchdog Status Register (WDOG_STAT)

The 32-bit Watchdog Status register (WDOG_STAT), shown in Figure 15-33, contains the current count value of the watchdog timer. Reads to WDOG_STAT return the current count value. When the watchdog timer is enabled, WDOG_STAT is decremented by 1 on each SCLK cycle. When WDOG_STAT reaches 0, the watchdog timer stops counting and the event selected in the Watchdog Control register (WDOG_CTL) is generated.

Values cannot be stored directly in WDOG_STAT, but are instead copied from WDOG_CNT. This can happen in two ways.

- While the watchdog timer is disabled, writing the WDOG_CNT register pre-loads the WDOG_STAT register.
- While the watchdog timer is enabled, writing the WDOG_STAT register loads it with the value in WDOG_CNT.

When the processor executes a write (of an arbitrary value) to WDOG_STAT, the value in WDOG_CNT is copied into WDOG_STAT. Typically, software sets the value of WDOG_CNT at initialization, then periodically writes to WDOG_STAT before the watchdog timer expires. This reloads the watchdog

timer with the value from WDOG_CNT and prevents generation of the selected event.

WDOG_STAT is a 32-bit unsigned system MMR that must be accessed with 32-bit reads and writes.

If the user does not reload the counter before SCLK*Count register cycles, a Watchdog interrupt or reset is generated and the WDRO bit in the Watchdog Control register is set. When this happens the counter will stop decrementing and will remain at zero.

If the counter is enabled with a zero loaded to the counter, the WDRO bit of the Watchdog Control register is set immediately and the counter remains at zero and does not decrement.

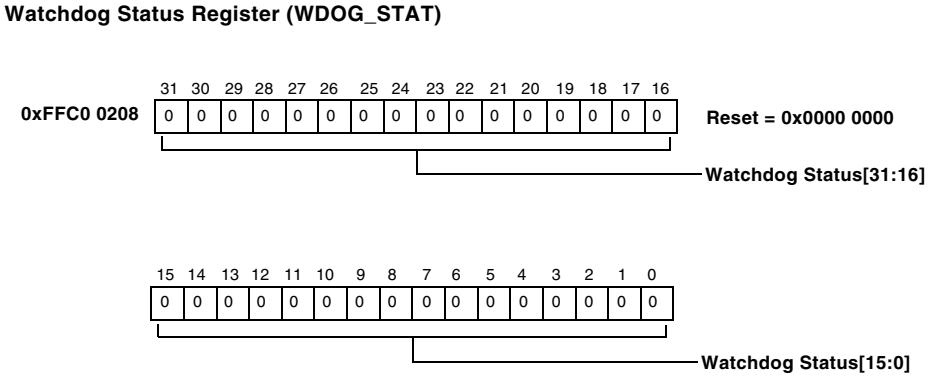

Figure 15-33. Watchdog Status Register

Watchdog Control Register (WDOG_CTL)

The Watchdog Control register (WDOG_CTL), shown in Figure 15-34, is a 16-bit system memory-mapped register used to control the watchdog timer.

The WDEV[1:0] field is used to select the event that is generated when the watchdog timer expires. Note that if the general-purpose interrupt option

is selected, the System Interrupt Mask register (SIC_IMASK) should be appropriately configured to unmask that interrupt. If the generation of watchdog events is disabled, the watchdog timer operates as described, except that no event is generated when the watchdog timer expires.

The WDEN[7:0] field is used to enable and disable the watchdog timer. Writing any value other than the disable value into this field enables the watchdog timer. This multibit disable key minimizes the chance of inadvertently disabling the watchdog timer.

Software can determine whether the timer has rolled over by interrogating the WDRO status bit of the Watchdog Control register. This is a sticky bit that is set whenever the watchdog timer count reaches 0 and cleared only by disabling the watchdog timer and then writing a 1 to the bit.

Watchdog Control Register (WDOG_CTL)

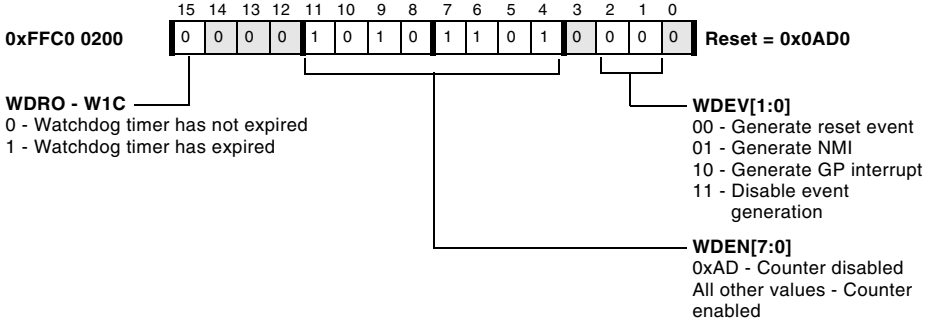

Figure 15-34. Watchdog Control Register

Note that when the processor is in Emulation mode, the watchdog timer counter will not decrement even if it is enabled.

16 EXTERNAL BUS INTERFACE UNIT

The External Bus Interface Unit (EBIU) provides glueless interfaces to external memories. The processor supports synchronous DRAM (SDRAM) and is compliant with the PC100 and PC133 SDRAM standards. The EBIU also supports asynchronous interfaces such as SRAM, ROM, FIFOs, flash memory, and ASIC/FPGA designs.

Overview

The EBIU services requests for external memory from the core or from a DMA channel. The priority of the requests is determined by the External Bus Controller. The address of the request determines whether the request is serviced by the EBIU SDRAM Controller or the EBIU Asynchronous Memory Controller.

The EBIU is clocked by the system clock (`SCLK`). All synchronous memories interfaced to the processor operate at the `SCLK` frequency. The ratio between core frequency and `SCLK` frequency is programmable using a phase-locked loop (PLL) system memory-mapped register (MMR). For more information, see “Core Clock/System Clock Ratio Control” on page 8-4.

The external memory space is shown in Figure 16-2 on page 16-4. One memory region is dedicated to SDRAM support. SDRAM interface timing and the size of the SDRAM region are programmable. The SDRAM memory space can range in size from 16 to 128 MB in each bank. The start address of the SDRAM memory space is 0x0000 0000. The area from

the end of the SDRAM memory space up to address 0x2000 0000 is reserved.

The next four regions (from 0x2000 0000 to 0x3000 0000) are dedicated to supporting asynchronous memories. Each asynchronous memory region is 64MByte and can be independently programmed to support different memory device characteristics. Each region has its own memory select output pin from the EBIU.

The next region is reserved memory space. References to this region do not generate external bus transactions. Writes have no effect on external memory values, and reads return undefined values. The EBIU generates an error response on the internal bus, which will generate a hardware exception for a core access or will optionally generate an interrupt from a DMA channel.

Block Diagram

Figure 16-1 is a conceptual block diagram of the EBIU and its interfaces. Signal names shown with an overbar are active low signals.

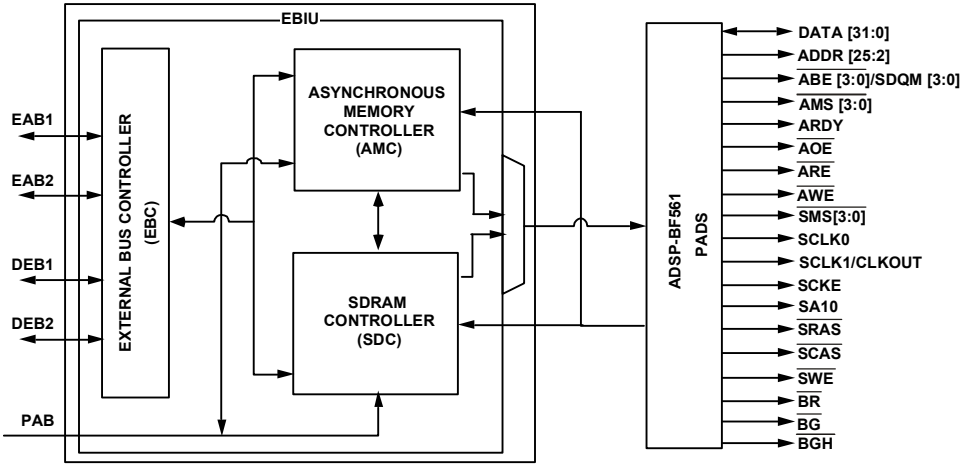

Figure 16-1. External Bus Interface Unit (EBIU)

Since only one external memory device can be accessed at a time, control, address, and data pins for each memory type are multiplexed together at the pins of the device. The Asynchronous Memory Controller (AMC) and the SDRAM Controller (SDC) effectively arbitrate for the shared pin resources.

Internal Memory Interfaces

The EBIU functions as a slave on three buses internal to the processor:

- External Access Bus (EAB), mastered by the core memory management unit on behalf of external bus requests from the core
- DMA External Bus (DEB), mastered by the DMA controller on behalf of external bus requests from any DMA channel
- Peripheral Access Bus (PAB), mastered by the core on behalf of system MMR requests from the core

These are synchronous interfaces, clocked by SCLK, as are the EBIU and Pads registers. The EAB provides access to both asynchronous external memory and synchronous DRAM external memory. The external access is controlled by either the Asynchronous Memory Controller (AMC) or the SDRAM Controller (SDC), depending on the internal address used to access the EBIU. Since the AMC and SDC share the same interface to the external pins, access is sequential and must be arbitrated based on requests from the EAB.

The third bus (PAB) is used only to access the memory-mapped control and status registers of the EBIU. The PAB connects separately to the AMC and SDC; it does not need to arbitrate with or take access cycles from the EAB bus.

The External Bus Controller (EBC) logic must arbitrate access requests for external memory coming from the EAB and DEB buses. The EBC logic routes read and write requests to the appropriate memory controller based

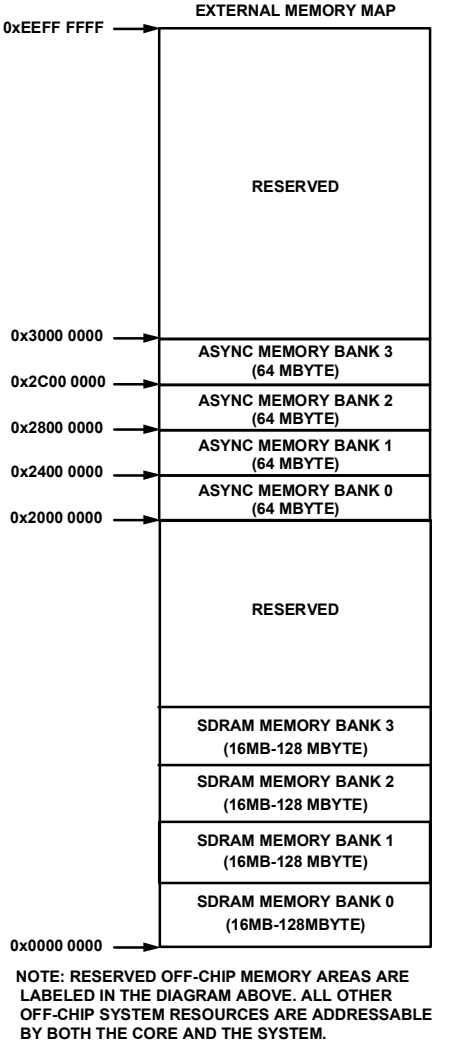

Figure 16-2. External Memory Map

on the bus selects. The AMC and SDC compete for access to the shared resources in the Pads logic. This competition is resolved in a pipelined fashion, in the order dictated by the EBC arbiter. Transactions from the core have priority over DMA accesses in most circumstances. However, if

the DMA controller detects an excessive backup of transactions, it can request its priority to be temporarily raised above the core.

External Memory Interfaces

Both the AMC and the SDC share the external interface address and data pins, as well as some of the control signals. These pins are shared:

- ADDR25–2, address bus
- DATA31–0, data bus
- $\overline{\text{ABE3–0}}$/SDQM3–0, AMC byte enables/SDC data masks
- $\overline{\text{BR}}$, $\overline{\text{BG}}$, $\overline{\text{BGH}}$, external bus access control signals

No other signals are multiplexed between the two controllers.

Table 16-1 and Table 16-2 describe the signals associated with each interface.

Table 16-1. Asynchronous Memory Interface Signals

Pad	Pin Type	Description
DATA31–0	I/O	External Data Bus
ADDR25–2	O	External Address Bus
$\overline{\text{AMS3–0}}$	O	Asynchronous Memory Selects
$\overline{\text{AWE}}$	O	Asynchronous Memory Write Enable
$\overline{\text{ARE}}$	O	Asynchronous Memory Read Enable
Pin Types: I = Input, O = Output		

Table 16-1. Asynchronous Memory Interface Signals (Cont'd)

Pad	Pin Type	Description
AOE	O	Asynchronous Memory Output Enable In most cases, the AOE pin should be connected to the OE pin of an external memory-mapped asynchronous device. Please refer to the applicable processor datasheet for specific timing information between the AOE and ARE signals to determine which interface signal should be used in your system.
ARDY	I	Asynchronous Memory Ready Response Note this is a synchronous input
ABE3–0/SDQM3–0	O	Byte Enables
SCLK0	O	System Clock Output
CLKOUT/SCLK1	O	System Clock Output
Pin Types: I = Input, O = Output		

Table 16-2. SDRAM Interface Signals

Pad	Pin Type	Description
DATA31–0	I/O	External Data Bus
ADDR19–18, ADDR16–2	O	External Address Bus Connect to SDRAM Address pins. Bank address is output on ADDR19–18 and should be connected to SDRAM BA1–0 pins.
SRAS	O	SDRAM Row Address Strobe pin Connect to SDRAM's RAS pin.
SCAS	O	SDRAM Column Address Strobe pin Connect to SDRAM's CAS pin.
SWE	O	SDRAM Write Enable pin Connect to SDRAM's WE pin.
ABE3–0/ SDQM3–0	O	SDRAM Data Mask pins Connect to SDRAM's DQM3–0 pins.
Pin Types: I = Input, O = Output		

Table 16-2. SDRAM Interface Signals (Cont'd)

Pad	Pin Type	Description
SMS3–0	O	Memory Select pin (of external memory bank configured for SDRAM) Connect to SDRAM's CS (Chip Select) pin. Active Low.
SA10	O	SDRAM A10 pin SDRAM interface uses this pin to be able to do refreshes while the AMC is using the bus. Connect to SDRAM's A10 pin.
SCKE	O	SDRAM Clock Enable pin Connect to SDRAM's SCKE pin.
SCLK0	O	SDRAM Clock Output pin 0 Switches at system clock frequency. Connect to the SDRAM's CLK pin.
CLKOUT/SCLK1	O	SDRAM Clock Output pin 1 Same frequency as SCLK0. Provided to reduce capacitive loading on SCLK0. Connect to the SDRAM CLK pin.
Pin Types: I = Input, O = Output		

EBIU Programming Model

This section describes the programming model of the EBIU. This model is based on system memory-mapped registers used to program the EBIU.

There are six control registers and one status register in the EBIU. They are:

- Asynchronous Memory Global Control register (EBIU_AMGCTL)
- Asynchronous Memory Bank Control 0 register (EBIU_AMBCTL0)
- Asynchronous Memory Bank Control 1 register (EBIU_AMBCTL1)
- SDRAM Memory Global Control register (EBIU_SDGCTL)
- SDRAM Memory Bank Control register (EBIU_SDBCTL)

- SDRAM Refresh Rate Control register (EBIU_SDRRC)
- SDRAM Control Status register (EBIU_SDSTAT)

Each of these registers is described in detail in the AMC and SDC sections later in this chapter.

Error Detection

The EBIU responds to any bus operation which addresses the range of 0x0000 0000 – 0xEEFF FFFF, even if that bus operation addresses reserved or disabled memory or functions. It responds by completing the bus operation (asserting the appropriate number of acknowledges as specified by the bus master) and by asserting the bus error signal for these error conditions:

- Any access to reserved off-chip memory space
- Any access to a disabled external memory bank
- Any access to an unpopulated area of an SDRAM memory bank

If the core requested the faulting bus operation, the bus error response from the EBIU is gated into the HWE interrupt internal to the core (this interrupt can be masked off in the core). If a DMA master requested the faulting bus operation, then the bus error is captured in that controller and can optionally generate an interrupt to the core.

Asynchronous Memory Interface

The asynchronous memory interface allows a glueless interface to a variety of memory and peripheral types. These include SRAM, ROM, EPROM, flash memory, and FPGA/ASIC designs. Four asynchronous memory regions are supported. Each has a unique memory select associated with it, shown in Table 16-3.

Table 16-3. Asynchronous Memory Bank Address Range

Memory Bank Select	Address Start	Address End
$\overline{\text{AMS3}}$	2C00 0000	2FFF FFFF
$\overline{\text{AMS2}}$	2800 0000	2BFF FFFF
$\overline{\text{AMS1}}$	2400 0000	27FF FFFF
$\overline{\text{AMS0}}$	2000 0000	23FF FFFF

Asynchronous Memory Address Decode

The address range allocated to each asynchronous memory bank is fixed at 64 MB; however, not all of an enabled memory bank need be populated. Unlike the SDRAM memory, which may need to support very large memory structures spanning multiple memory banks, it should be relatively easy to constrain code and data structures to fit within one of the supported asynchronous memory banks, because of the nature of the types of code or data that is stored here.

Note accesses to unpopulated memory of partially populated AMC banks do not result in a bus error and will alias to valid AMC addresses.

The asynchronous memory signals are defined in Table 16-1. The timing of these pins is programmable to allow a flexible interface to devices of different speeds. For example interfaces, see Chapter 17, “System Design”.

Asynchronous Memory Global Control Register (EBIU_AMGCTL)

The Asynchronous Memory Global Control register configures global aspects of the controller. It contains bank enables and other information as described in this section. This register should not be programmed while the AMC is in use. The `EBIU_AMGCTL` register should be the last control register written to when configuring the processor to access external

memory-mapped asynchronous devices. Writes to this register should be followed by an SSYNC instruction to ensure that the asynchronous bank is enabled before an access is attempted.

Asynchronous Memory Global Control (EBIU_AMGCTL)

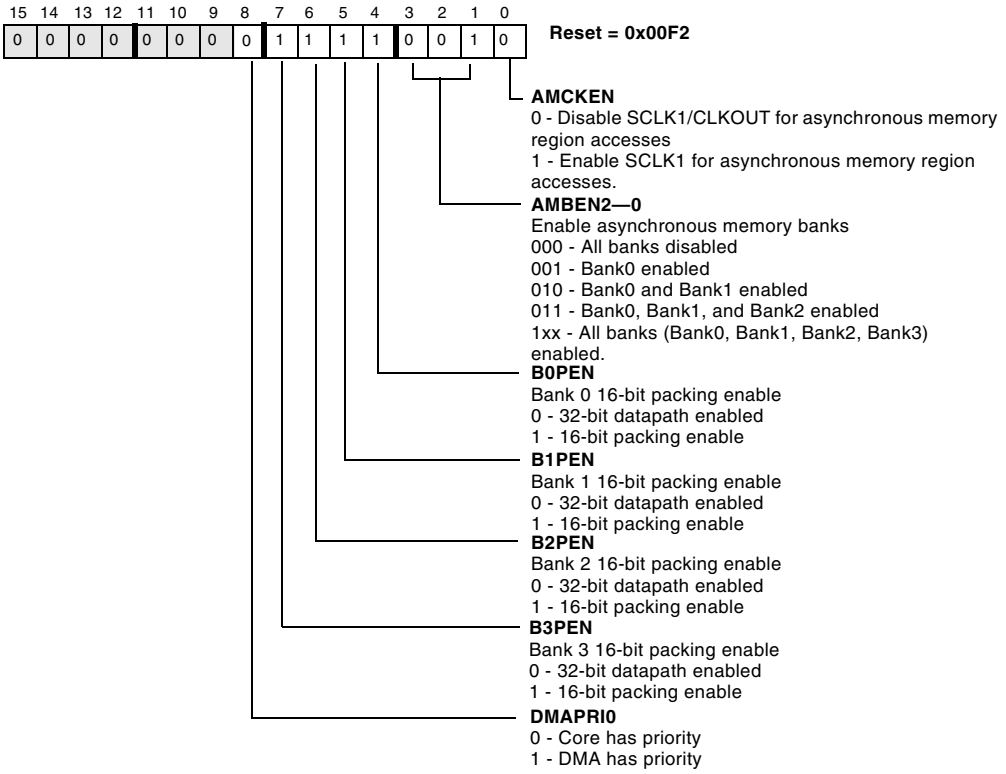

Figure 16-3. Asynchronous Memory Global Control (EBIU_AMGCTL)

Within the EBIU_AMGCTL register, the AMCKEN bit is used for external devices that are connected to CLKOUT. When it is needed, the CLKOUT/SCLK1 pin can be enabled by setting the AMCKEN bit in the EBIU_AMGCTL register. In systems that do not use CLKOUT/SCLK1, set the AMCKEN bit to 0.

Writing a “1” to the AMBEN2-0 bits will enable specific asynchronous memory banks, while a writing a “0” will disable specific banks. If a bus operation accesses a disabled asynchronous memory bank, the EBIU responds by acknowledging the transfer and asserting the error signal on the requesting bus. The error signal propagates back to the requesting bus master. This generates a hardware exception to the core, if it is the requester. For DMA mastered requests, the error is captured in the respective status register. If a bank is not fully populated with memory, then the memory likely aliases into multiple address regions within the bank. This aliasing condition is not detected by the EBIU, and no error response is asserted.

The 16-bit packing enable bits (B0PEN, B1PEN, B2PEN and B3PEN) are available to control the access size for each asynchronous memory bank separately. When one of these bits is set to one, any request to the corresponding AMC (Asynchronous Memory Controller) bank will result in a 16-bit access to external devices. Regardless of whether the request comes from a 8-bit, 16-bit, 32-bit data fetch or whether the fetch originates from the DMA controller, the EBIU will perform a 16-bit access. In the case of 32-bit requests to the AMC, a dual 16-bit access will be made to fetch 32-bits.

For more details on the use of $\overline{\text{ABE3-0}}$, please refer to “Byte Enables” on page 16-22.

When a 16-bit packing enable bit is set to zero (that is, “cleared”), the AMC treats any request to the dedicated bank as a 32-bit access to external device. Regardless of the kind of access (8,16 or 32-bit fetch), the external bus will perform a 32-bit access.

The CDPRIO bit configures the AMC to control the priority requests that occur simultaneously to the EBIU from either the processor core or the DMA controller. When this bit is set to zero, requests from either core have priority over a request from the DMA controller to the AMC, unless the DMA is “urgent”. When the CDPRIO bit is set, all requests from the DMA controllers, including the memory DMAs, have priority over core

accesses. For the purposes of this discussion, core accesses include both data fetches and instruction fetches.

Asynchronous Memory Bank Control Registers (EBIU_AMBCTL0, EBIU_AMBCTL1)

The EBIU asynchronous memory controller has two memory bank control registers (Figure 16-4 and Figure 16-5). They contain bits for counters for setup, strobe, and hold time; bits to determine memory type and size; and bits to configure use of ARDY. These registers should not be programmed while the AMC is in use.

The timing characteristics of the AMC can be programmed using these four parameters:

- Setup: the time between the beginning of a memory cycle ($\overline{\text{AMSx}}$ low) and the read-enable assertion ($\overline{\text{ARE}}$ low) or write-enable assertion ($\overline{\text{AWE}}$ low)
- Read Access: the time between read-enable assertion ($\overline{\text{ARE}}$ low) and de-assertion ($\overline{\text{ARE}}$ high)
- Write Access: the time between write-enable assertion ($\overline{\text{AWE}}$ low) and de-assertion ($\overline{\text{AWE}}$ high)
- Hold: the time between read-enable de-assertion ($\overline{\text{ARE}}$ high) or write-enable de-assertion ($\overline{\text{AWE}}$ high) and the end of the memory cycle ($\overline{\text{AMSx}}$ high)

Each of these parameters can be programmed in terms of EBIU clock cycles. In addition, there are minimum values for these parameters:

- Setup ≥ 1 cycle
- Read Access ≥ 1 cycle

- Write Access ≥ 1 cycle
- Hold ≥ 0 cycles

Avoiding Bus Contention

Because the three-stated data bus is shared by multiple devices in a system, be careful to avoid contention. Contention causes excessive power dissipation and can lead to device failure. Contention occurs during the time one device is getting off the bus and another is getting on. If the first device is slow to three-state and the second device is quick to drive, the devices contend.

There are two cases where contention can occur. The first case is a read followed by a write to the same memory space. In this case, the data bus drivers can potentially contend with those of the memory device addressed by the read. The second case is back-to-back reads from two different memory spaces. In this case, the two memory devices addressed by the two reads could potentially contend at the transition between the two read operations.

To avoid contention, program the turnaround time (Bank Transition Time) appropriately in the Asynchronous Memory Bank Control registers. This feature allows software to set the number of clock cycles between these types of accesses on a bank-by-bank basis. Minimally, the EBIU provides one cycle for the transition to occur.

ARDY Input Control

Each bank can be programmed to sample the ARDY input after the read or write access timer has counted down or to ignore this input signal. If enabled and disabled at the sample window, ARDY can be used to extend

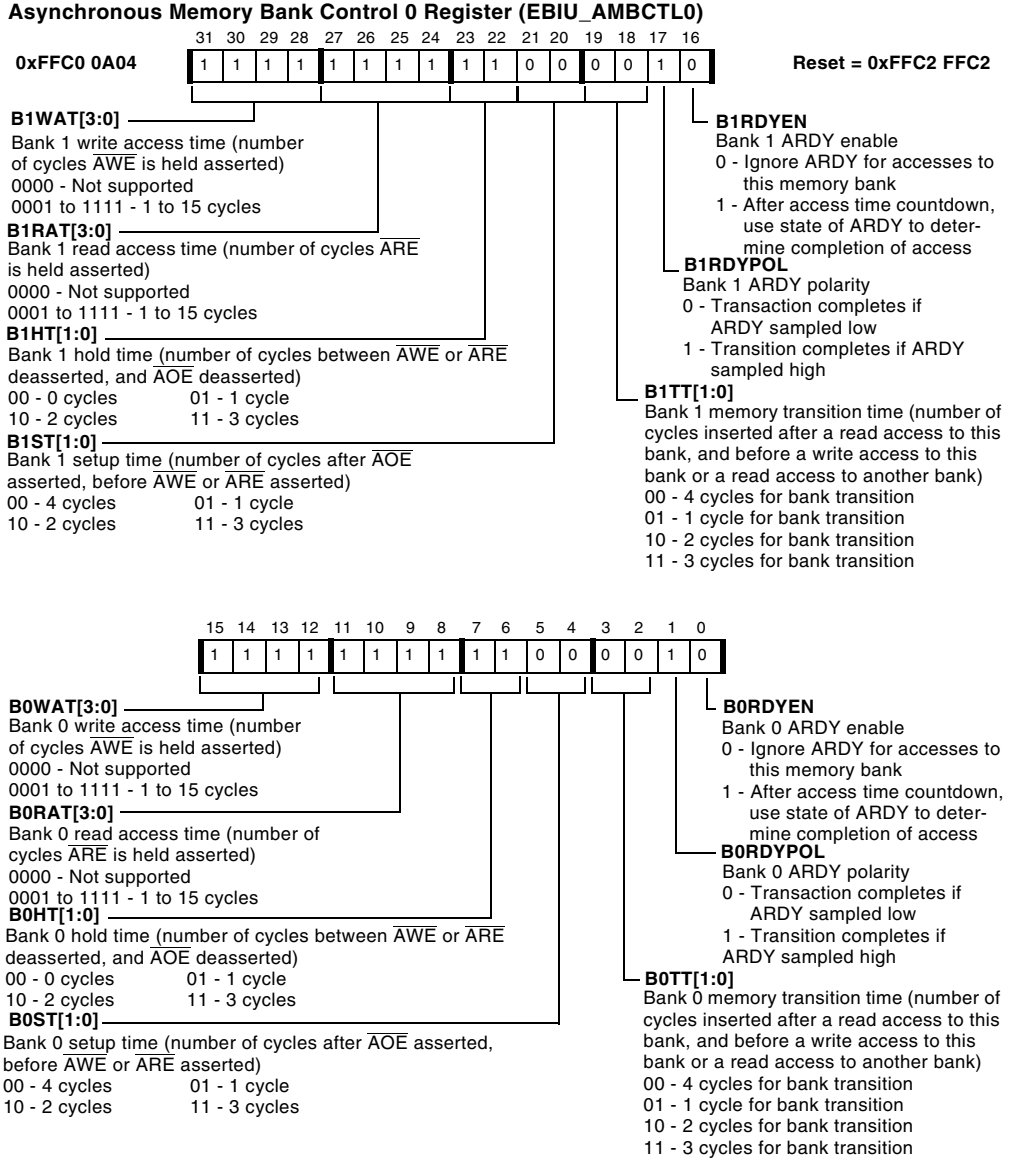

Figure 16-4. Asynchronous Memory Bank Control 0 Register

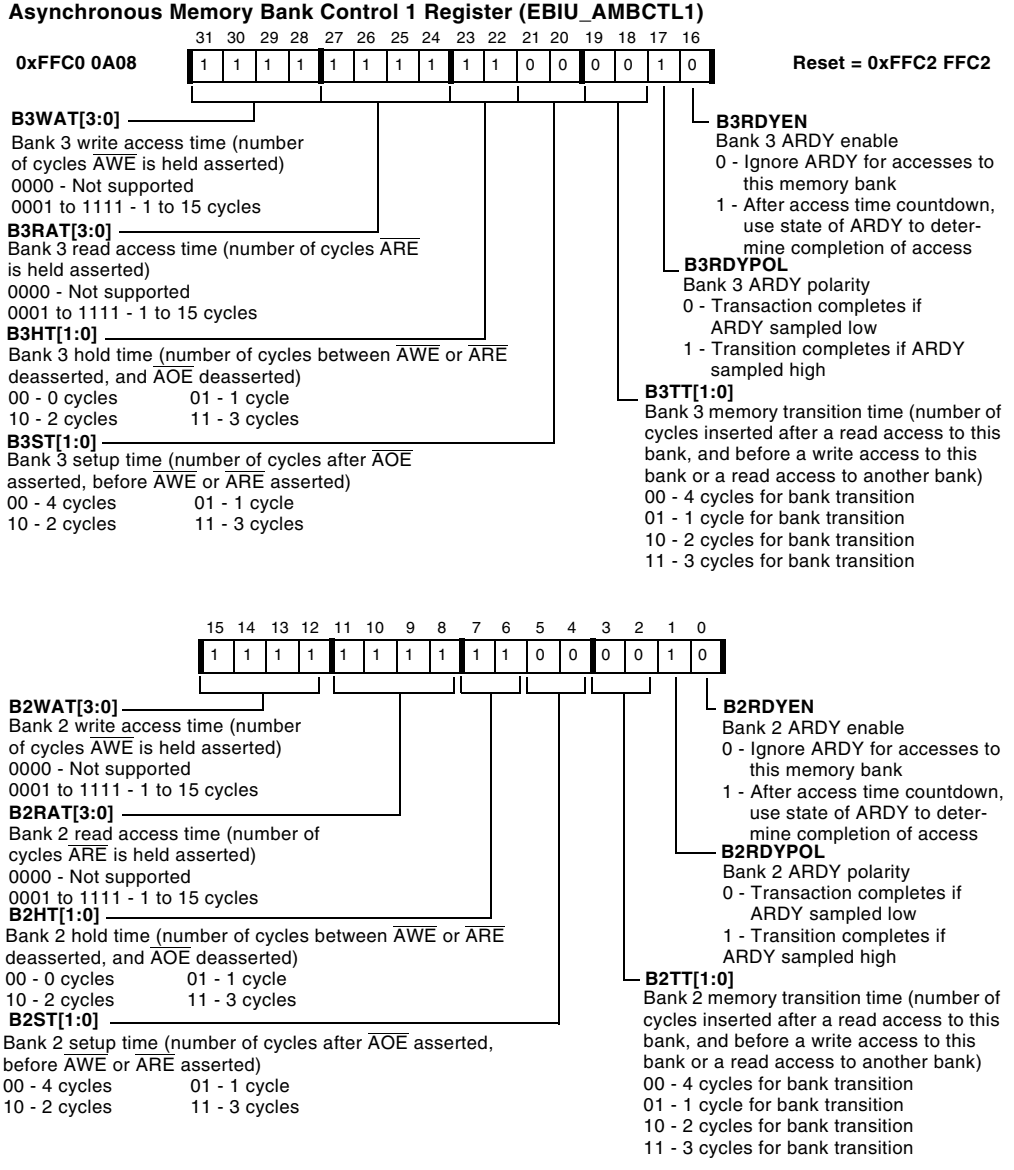

Figure 16-5. Asynchronous Memory Bank Control 1 Register

the access time as required. Note ARDY is synchronously sampled, therefore:

- Assertion and de-assertion of ARDY to the processor must meet the *ADSP-BF561 Embedded Processor Data Sheet* setup and hold times. Failure to meet these synchronous specifications could result in metastable behavior internally. The processor's CLKOUT signal should be used to ensure synchronous transitions of ARDY.
- The ARDY pin must be stable (either asserted or deasserted) at the external interface on the cycle before the internal bank counter reaches 0; that is, more than one CLKOUT/SCLK1 cycle before the scheduled rising edge of $\overline{\text{AWE}}$ or $\overline{\text{ARE}}$. This will determine whether the access is extended or not.
- Once the transaction has been extended by the assertion of ARDY, the transaction completes in the cycle after ARDY is sampled as asserted.

The polarity of ARDY is programmable on a per-bank basis. Since ARDY is not sampled until an access is in progress to a bank in which the ARDY enable is asserted, ARDY does not need to be driven by default. For more information, see "Adding Additional Wait States" on page 16-19.

Programmable Timing Characteristics

Asynchronous Reads

Figure 16-6 shows an asynchronous-read bus cycle with timing programmed as setup = 2 cycles, read access = 2 cycles, hold = 1 cycle, and transition time = 1 cycle.

Asynchronous read bus cycles proceed as follows:

1. At the start of the setup period, the $\overline{\text{AMSx}}$ and $\overline{\text{AOE}}$ assert. The address bus becomes valid. $\overline{\text{ABE3–0}}$ are low during the read.
2. At the beginning of the read access period and after the 2 setup cycles, $\overline{\text{ARE}}$ asserts.
3. At the beginning of the hold period, read data is sampled on the rising edge of the EBIU clock. The $\overline{\text{ARE}}$ pin deasserts after this rising edge.
4. At the end of the hold period, $\overline{\text{AOE}}$ deasserts unless this bus cycle is followed by another asynchronous read to the same memory space. Also $\overline{\text{AMSx}}$ deasserts unless the next cycle is to the same memory bank.
5. Unless another read of the same memory bank is queued internally, the AMC appends the programmed number of memory transition time cycles.

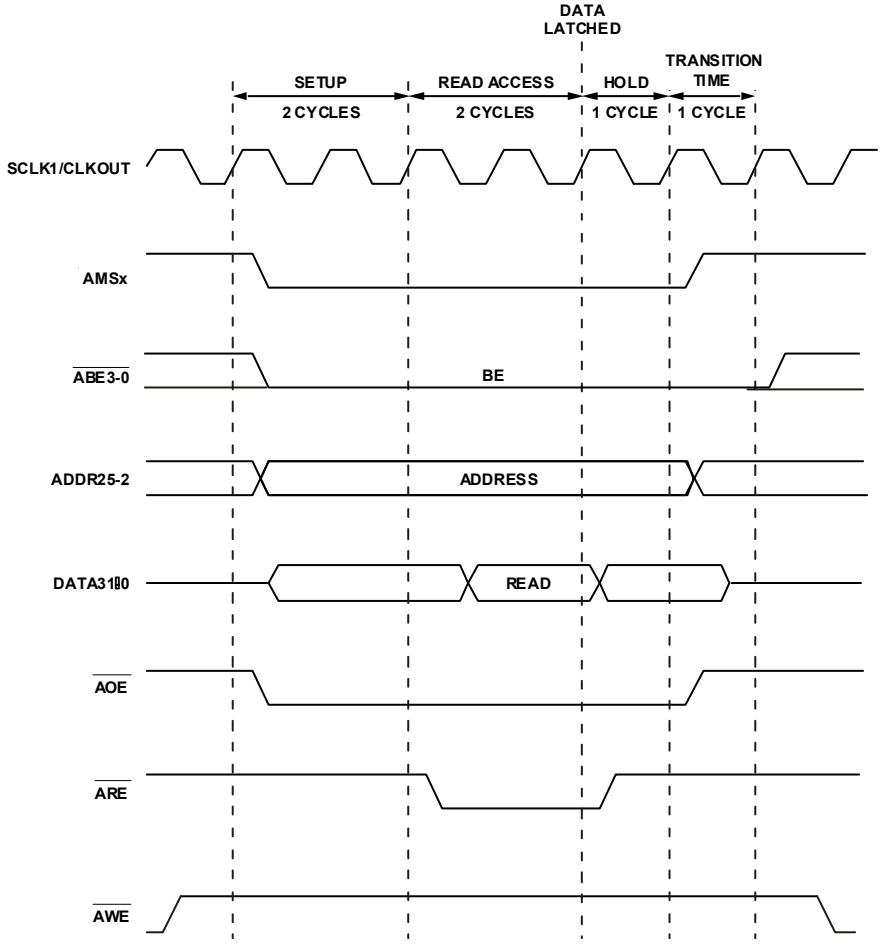

Figure 16-6. Asynchronous Read Bus Cycles

Asynchronous Writes

Figure 16-7 shows an asynchronous write bus cycle followed by an asynchronous read cycle to the same bank, with timing programmed as

setup = 2 cycles, write access = 2 cycles, read access = 3 cycles, hold = 1 cycle, and transition time = 1 cycle.

Asynchronous write bus cycles proceed as follows:

1. At the start of the setup period, the $\overline{\text{AMSx}}$, the address bus, data buses, and $\overline{\text{ABE3–0}}$ become valid. For more details on $\overline{\text{ABE3-0}}$, refer to "Byte Enables" on page 16-22.
2. At the beginning of the write access period, $\overline{\text{AWE}}$ asserts.
3. At the beginning of the hold period, $\overline{\text{AWE}}$ deasserts.

Asynchronous read bus cycles proceed as follows:

1. At the start of the setup period, the $\overline{\text{AMSx}}$ and $\overline{\text{AOE}}$ asserts. The address bus becomes valid. $\overline{\text{ABE3–0}}$ are low during the read.
2. At the beginning of the read access period, $\overline{\text{ARE}}$ asserts.
3. At the beginning of the hold period, read data is sampled on the rising edge of the EBIU clock. $\overline{\text{ARE}}$ deasserts after this rising edge.
4. At the end of the hold period, $\overline{\text{AOE}}$ deasserts unless this bus cycle is followed by another asynchronous read to the same memory space. Also $\overline{\text{AMSx}}$ deasserts unless the next cycle is to the same memory bank.
5. Unless another read of the same memory bank is queued internally, the AMC appends the programmed number of memory transition time cycles.

Adding Additional Wait States

The ARDY pin is used to insert extra wait states. The input is sampled synchronously with the EBIU internal clock. The EBIU starts sampling ARDY on the clock cycle before the end of the programmed strobe period. If ARDY is sampled deasserted, the access period is extended. The ARDY pin is

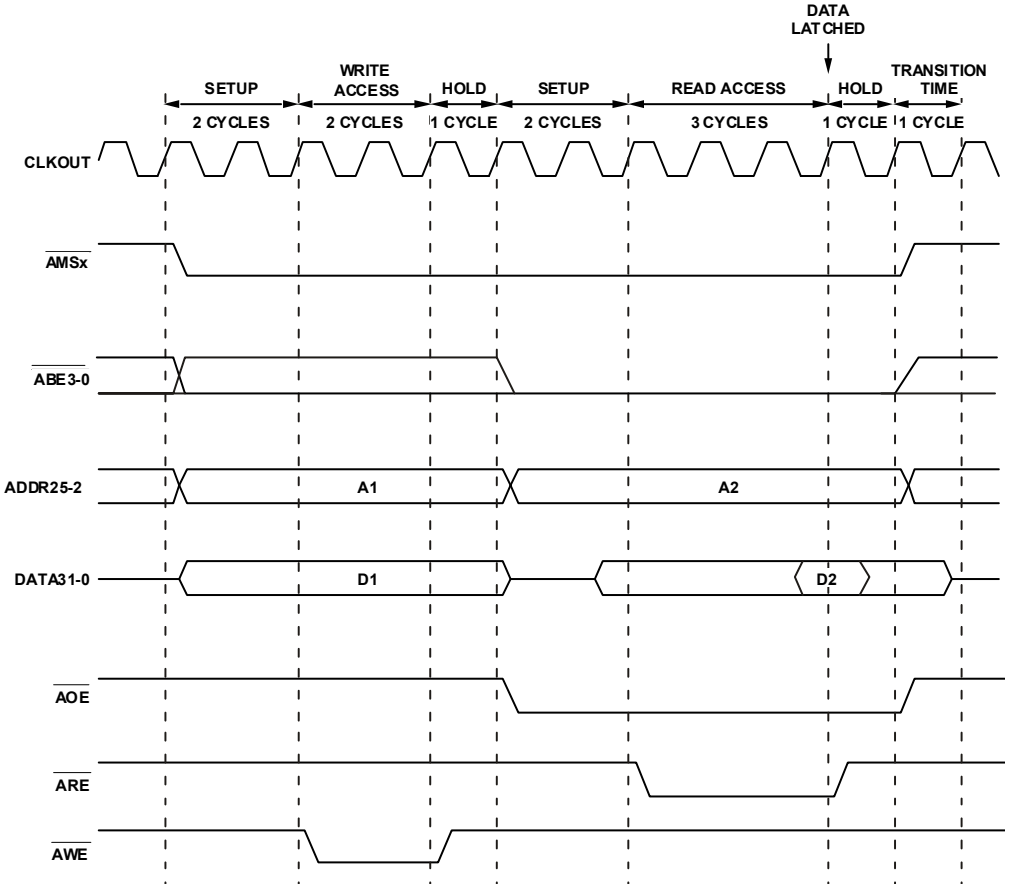

Figure 16-7. Asynchronous Write and Read Bus Cycles

then sampled on each subsequent clock edge. Read data is latched on the clock edge after ARDY is sampled asserted. The read- or write-enable remains asserted for one clock cycle after ARDY is sampled asserted. An example of this behavior is shown in Figure 16-8, where setup = 2 cycles, read access = 4 cycles, and hold = 1 cycle. Note the read access period

must be programmed to a minimum of two cycles to make use of the ARDY input.

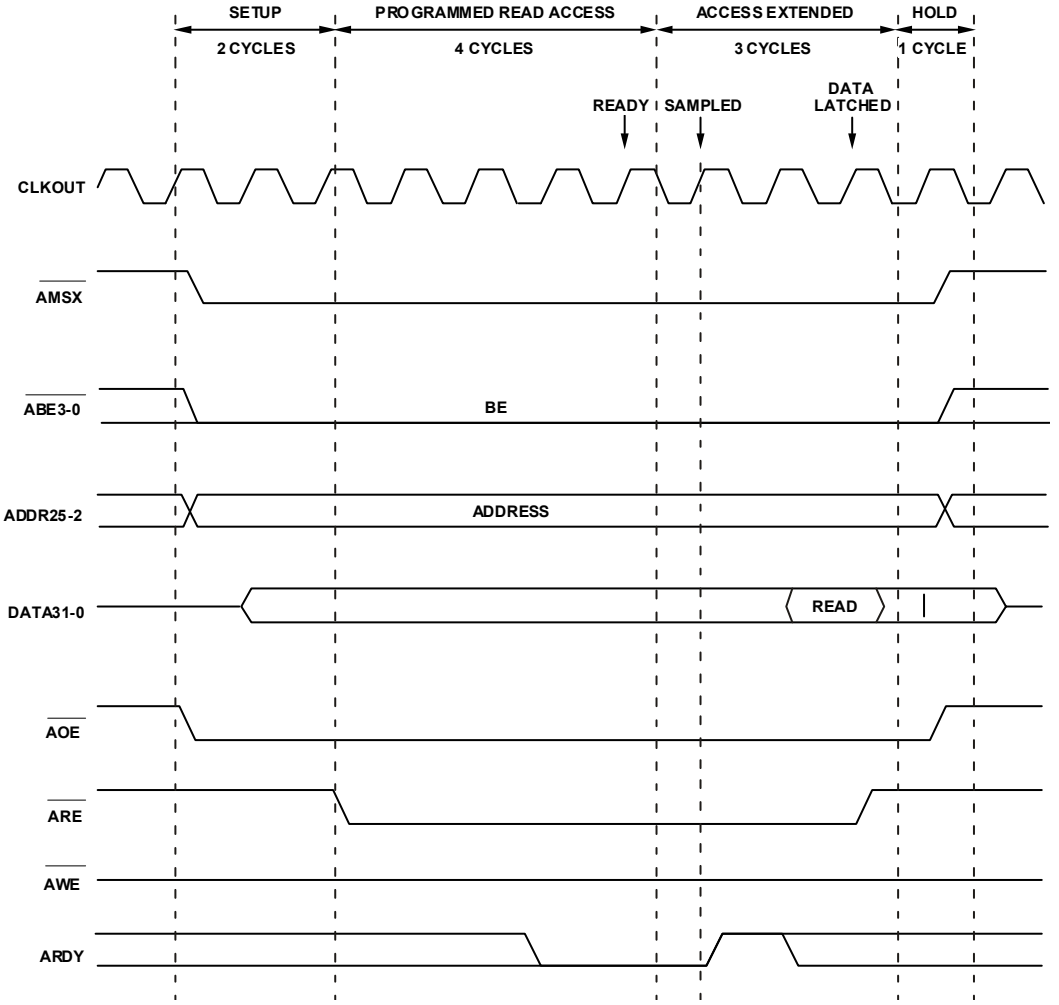

Figure 16-8. Inserting Wait States Using ARDY

Byte Enables

The EBIU provides byte enable pins to allow Blackfin to do efficient byte wide arithmetic and byte wide processing in external memory. Note that any read access to external devices will not require extra byte enable pins, therefore $\overline{\text{ABE[3:0]}}$ remain low during read accesses. In general, there are two ways to modify a single byte or even a 16-bit word within the 32-bit interface. First, it can be realized by a read/modify/write sequence. However, this is not very efficient because multiple accesses are required (for example, it takes many cycles for reads and writes to external memory). A second option is to modify just a specific byte for a 16-bit or 32-bit device using the $\overline{\text{ABE[3:0]}}$ pins.

The $\overline{\text{ABE[3]}}$ pin has two different functions. When the AMC is configured to do 16-bit data packing via the Asynchronous Memory Global Control Register, the $\overline{\text{ABE[3]}}$ pin functions as the least significant bit of the address bus ($\overline{\text{ABE[3]}}$ = A1). Since the EBIU does not support byte addressing, the address pin A0 does not exist. Either $\overline{\text{ABE[3]}}$ (16-bit data packing mode) or A2 (32-bit mode) has control over the least significant address bit. Nevertheless, 8-bit devices can be accessed by using 16-bit accesses and taking the lower byte and ignoring the upper byte.

When the AMC is configured to access 32-bit devices (16-bit data packing is disabled) $\overline{\text{ABE[3:0]}}$ is part of the four byte select lines, which make byte/word modification possible.

Accessing 8-Bit/16-Bit Devices (16-Bit Data Packing is Enabled)

For any access to external devices, none of the pins ($\overline{\text{ABE[2]}}$, $\overline{\text{ABE1}}$ or $\overline{\text{ABE0}}$) need to be used (remember $\overline{\text{ABE3}}$ is taken as A1). Nevertheless, $\overline{\text{ABE1}}$ and $\overline{\text{ABE0}}$ can be used for byte wide write accesses to devices that support byte accesses (Chapter 17, “System Design”).

Since two sequential 8-bit requests result in two accesses to the same 16-bit address, it is recommended that eight bit devices be accessed by using two 16-bit requests, simply keeping the lower byte of each access.

An example of 16-bit memories supporting byte accesses can be found in Chapter 17, "System Design". In this case, $\overline{\text{ABE0}}$ and $\overline{\text{ABE1}}$ control the write access to lower and upper bytes.

Table 16-4. 8-Bit Write Accesses (16-Bit Memory)

	Byte 0	Byte 1
$\overline{\text{ABE0}}$	0	1
$\overline{\text{ABE1}}$	1	0

Accessing 32-Bit Devices (16-Bit Data Packing is Disabled)

In this mode, none of the pins ($\overline{\text{ABE3}}$, $\overline{\text{ABE2}}$, $\overline{\text{ABE1}}$, or $\overline{\text{ABE0}}$) are required to be connected. But they can be used to control byte or even 16-bit word write accesses. To perform a 16-bit write to external devices, $\overline{\text{ABE1}}$ or $\overline{\text{ABE0}}$ control accesses to the lower word, while $\overline{\text{ABE3}}$ and $\overline{\text{ABE2}}$ control accesses to the upper word. When performing 8-bit write accesses to external devices, $\overline{\text{ABE3}}$ controls write access to the upper byte and $\overline{\text{ABE2}}$ controls write accesses to the lower byte of the upper 16-bits. $\overline{\text{ABE1}}$ controls write accesses to the upper byte of the lower 16-bits and $\overline{\text{ABE0}}$ controls write accesses to the lower byte of the lower 16-bits.

Since two sequential 16-bit requests to the AMC result in a dual access to the same 32-bit address, it is recommended that 16-bit devices be accessed with 32-bit requests, simply keeping the lower word in this mode.

Table 16-5 and Table 16-6 show the state of the pins at different accesses.

Table 16-5. 16-Bit Write Accesses (32-Bit Memory)

	Lower 16-Bit Word	Upper 16-Bit Word
$\overline{ABE0}$	0	0
$\overline{ABE1}$	0	0
$\overline{ABE2}$	1	1
$\overline{ABE3}$	1	1

Table 16-6. 8-Bit Write Accesses (32-Bit Memory)

	Byte 0	Byte 1	Byte 2	Byte 3
$\overline{ABE0}$	0	1	1	1
$\overline{ABE1}$	1	0	1	1
$\overline{ABE2}$	1	1	0	1
$\overline{ABE3}$	1	1	1	0

SDRAM Controller (SDC)

Introduction

The SDRAM Controller (SDC) enables the processor to transfer data to and from Synchronous DRAM (SDRAM) with a maximum frequency specified in the *ADSP-BF561 Embedded Processor Data Sheet*. The processor supports a glueless interface with four external banks (controlled by the SMS3–0 Memory Select pins) of standard SDRAMs of 64 Mbit - 512 Mbit with configurations x4, x8, x16, and x32 as shown in the following tables, up to a maximum total capacity of 512 MBytes (MB) of SDRAM. The interface includes timing options to support additional buffers between the processor and SDRAM, to handle the capacitive loads of large memory arrays.

All inputs are sampled and all outputs are valid on the rising edge of the SDRAM clock output CLKOUT.

The EBIU SDC provides a glueless interface with standard SDRAMs. The SDRAM controller:

- Supports SDRAMs of 64 Mbit, 128 Mbit, 256 Mbit, and 512 Mbit with configurations of x4, x8, x16, and x32
- Supports up to 512 MB of SDRAM (128 MB in each bank)
- Each external bank is independently programmable in size
- Supports SDRAM page sizes of 1 kbyte, 2 kbyte, 4 kbyte, and 8 kbyte for 32-bit wide SDRAM banks and 512 byte, 1 kbyte, 2 kbyte, and 4 kbyte for 16-bit wide SDRAM banks.
- Supports four internal banks within the SDRAMs
- Supports interleaved SDRAM bank accesses
- Uses a programmable refresh counter to coordinate between varying clock frequencies and the SDRAM's required refresh rate
- Provides multiple timing options to support additional buffers between the processor and SDRAM
- Uses a separate pin (SA10) that enables the SDC to precharge SDRAM before issuing an Auto-Refresh or Self-Refresh command while the asynchronous memory controller has control of the EBIU port
- Supports self-refresh for standard and mobile SDRAMs
- Provides two SDRAM powerup options

Definition of Terms

The following are definitions used in the remainder of this chapter.

Bank Activate Command

The Bank Activate command causes the SDRAM to open an internal bank (specified by the bank address) in a row (specified by the row address). When the Bank Activate command is issued to the SDRAM, the SDRAM opens a new row address in the dedicated bank. The memory in the open internal bank and row is referred to as the open page. The Bank Activate command must be applied before a read or write command.

Burst Length

The burst length determines the number of words that the SDRAM device stores or delivers after detecting a single write or read command, respectively. The burst length is selected by writing certain bits in the SDRAM's Mode register during the SDRAM powerup sequence.

Although the SDC supports only Burst Length = 1 mode, during a burst to SDRAM, the SDC applies the read or write command every cycle and keeps accessing the data. Setting Burst Length = 1 does not reduce the performance throughput.

Burst Stop Command

The Burst Stop command is one of several ways to terminate or interrupt a burst read or write operation.

Since the SDRAM burst length is always hardwired to be 1, the SDC does not support the Burst Stop command.

Burst Type

The burst type determines the address order in which the SDRAM delivers burst data after detecting a read command or stores burst data after detecting a write command. The burst type is programmed in the SDRAM during the SDRAM powerup sequence.

Since the SDRAM burst length is always programmed to be 1, the burst type does not matter. However, the SDC always sets the burst type to sequential-accesses-only during the SDRAM powerup sequence.

CAS Latency (CL)

The Column Address Strobe (CAS) latency is the delay in clock cycles between when the SDRAM detects the read command and when it provides the data at its output pins. The CAS latency is programmed in the SDRAM Mode register during the powerup sequence.

The speed grade of the SDRAM and the application's clock frequency determine the value of the CAS latency. The SDC can support CAS latency of two or three clock cycles. The selected CAS latency value must be programmed into the SDRAM Memory Global Control register (EBIU_SDGCTL) before the SDRAM powerup sequence. See "SDRAM Memory Global Control Register (EBIU_SDGCTL)" on page 16-34.

CBR (CAS Before RAS) Refresh or Auto-Refresh

When the SDC refresh counter times out, the SDC precharges all four banks of SDRAM and then issues an Auto-Refresh command to them. This causes the SDRAMs to generate an internal CBR refresh cycle. When the internal refresh completes, all four internal SDRAM banks are precharged.

DQM3—0 Data I/O Mask Function

The SDQM3–0 pins provide a byte-masking capability on 8- or 16-bit writes to SDRAM. The DQM3–0 pins are used to block the output buffer of the SDRAM during partial write operations. The SDQM3–0 pins are not used to mask data on read cycles.

Internal Bank

There are several internal memory banks on a given SDRAM. The SDC supports interleaved accesses among the internal banks. The bank address can be thought of as part of the row address. The SDC assumes that all SDRAMs to which it interfaces have four internal banks and allows each activated bank to have a unique row address.

Mode Register

SDRAM devices contain an internal configuration register which allows specification of the SDRAM device's functionality. After powerup and before executing a read or write to the SDRAM memory space, the application must trigger the SDC to write the SDRAM's Mode register. The write of the SDRAM's Mode register is triggered by writing a 1 to the PSSE bit in the SDRAM Memory Global Control register (EBIU_SDGCTL) and then issuing a read or write transfer to the SDRAM address space. The initial read or write triggers the SDRAM powerup sequence to be run, which programs the SDRAM's Mode register with the CAS latency from the EBIU_SDGCTL register. This initial read or write to SDRAM takes many cycles to complete. Note for most applications, the SDRAM powerup sequence and writing of the Mode register needs to be done only once. Once the powerup sequence has completed, the PSSE bit should not be set again unless a change to the Mode register is desired.

Low power SDRAM devices may also contain an Extended Mode register. The EBIU enables programming of the Extended Mode register during powerup via the EMREN bit in the EBIU_SDGCTL register.

If the SDRAM mode register's CAS latency must be reprogrammed (subsequent powerup), the SDRAM should first be programmed in self-refresh mode.

Page Size

Page size is the amount of memory which has the same row address and can be accessed with successive read or write commands without needing to activate another row. The page size can be calculated for 32/16-bit SDRAM banks with the following formula:

- 32-bit SDRAM banks: page size = $2^{(CAW + 2)}$

where CAW is the column address width of the SDRAM, plus 2 because the SDRAM bank is 32 bits wide (2 address bit = 4 bytes).

- 16-bit SDRAM banks: page size = $2^{(CAW + 1)}$)

where CAW is the column address width of the SDRAM, plus 1 because the SDRAM bank is 16 bits wide (1 address bit = 2 bytes).

Precharge Command

The Precharge command closes a specific internal bank in the active page or all internal banks in the page.

SDRAM Bank

The SDRAM bank is a region of memory that can be configured to 16 MB, 32 MB, 64 MB, or 128 MB and is selected by the $\overline{\text{SMS3–0}}$ pin.

Do not confuse the “SDRAM internal banks” which are internal to the SDRAM and are selected with the bank address, with the “SDRAM bank” or “external bank” that is enabled by the $\overline{\text{SMS3–0}}$ pin.

Self-Refresh

When the SDRAM is in Self-Refresh mode, the SDRAM’s internal timer initiates Auto-Refresh cycles periodically, without external control input. The SDC must issue a series of commands including the Self-Refresh

command to put the SDRAM into this low power mode, and it must issue another series of commands to exit Self-Refresh mode. Entering Self-Refresh mode is programmable in the SDRAM Memory Global Control register (EBIU_SDGCTL) and any access to the SDRAM address space causes the SDC to exit the SDRAM from Self-Refresh mode. See "Entering and Exiting Self-Refresh Mode (SRFS)" on page 16-40.

t_{RAS}

This is the required delay between issuing a Bank Activate command and issuing a Precharge command, and between the Self-Refresh command and the exit from Self-Refresh. The TRAS bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is 4 bits wide and can be programmed to be 1 to 15 clock cycles long. See "Selecting the Bank Activate Command Delay (TRAS)" on page 16-43.

t_{RP}

This is the required delay between issuing a Precharge command and:

- issuing a Bank Activate command
- issuing an Auto-Refresh command
- issuing a Self-Refresh command

The TRP bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is three bits wide and can be programmed to be 1 to 7 clock cycles long. See "Selecting the Precharge Delay (TRP)" on page 16-44.

t_{RCD}

This is the required delay between a Bank Activate command and the start of the first Read or Write command. The TRCD bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is three bits wide and can be programmed to be from 1 to 7 clock cycles long.

t_{WR}

This is the required delay between a Write command (driving write data) and a Precharge command. The TWR bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is two bits wide and can be programmed to be from 1 to 3 clock cycles long.

t_{RC}

This is the required delay between issuing successive Bank Activate commands to the same SDRAM internal bank. This delay is not directly programmable. The t_{RC} delay must be satisfied by programming the TRAS and TRP fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

t_{RFC}

This is the required delay between issuing an Auto-Refresh command and a Bank Activate command and between issuing successive Auto-Refresh commands. This delay is not directly programmable and is assumed to be equal to t_{RC}. The t_{RC} delay must be satisfied by programming the TRAS and TRP fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

t_{XSR}

This is the required delay between exiting Self-Refresh mode and issuing the Auto-Refresh command. This delay is not directly programmable and is assumed to be equal to t_{RC}. The t_{RC} delay must be satisfied by programming the t_{RAS} and t_{RP} fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

t_{RRD}

This is the required delay between issuing a Bank A Activate command and a Bank B Activate command. This delay is not directly programmable and is assumed to be $t_{RCD} + 1$.

Example SDRAM System Block Diagrams

Figure 16-9 shows a block diagram of the SDRAM interface. In this example, the SDRAM interface connects to two 64 Mbit (x8) SDRAM devices to form one external bank of 16 Mbytes of memory. The same address and control bus feeds both SDRAM devices.

Figure 16-10 shows a block diagram of the ADSP-BF561 processor's SDRAM interface. In this example, the SDRAM interface connects to four 64 Mbit (x8) SDRAM devices, to form one external bank of 32 Mbytes of memory. The same address and control bus feeds all four SDRAM devices.

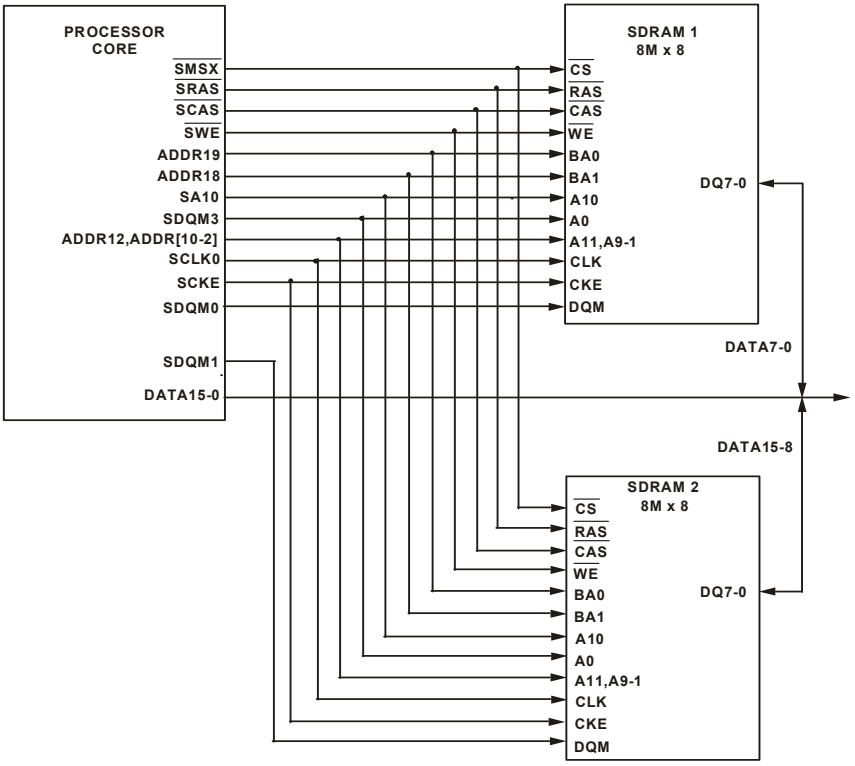

Figure 16-9. 16 MB SDRAM System Interface Example

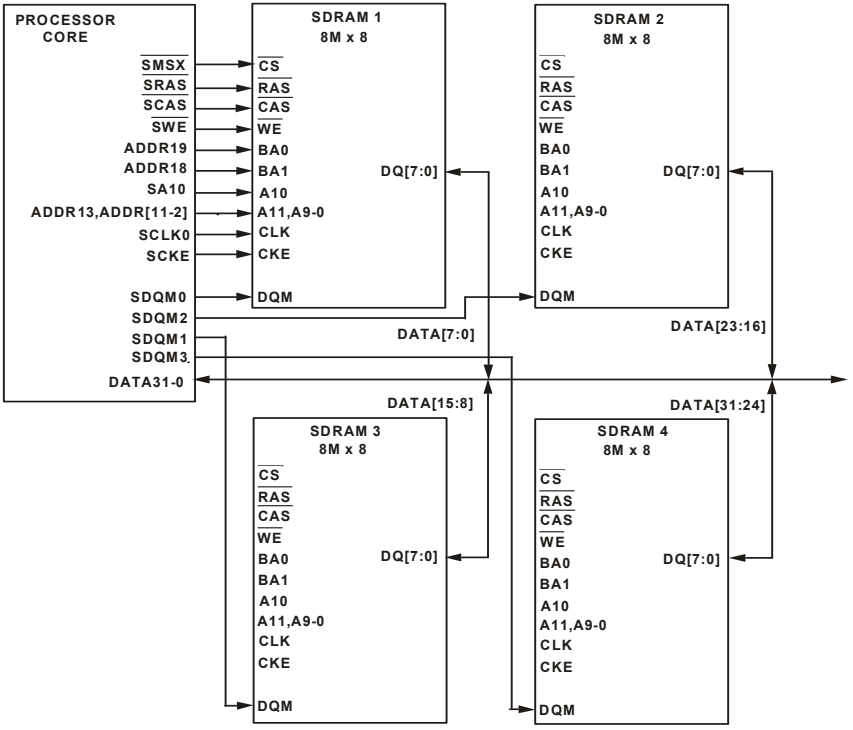

Figure 16-10. 32 MB SDRAM System Interface Example

SDQM Operation

The SDQM3–0 (Data I/O Mask) pins enable the SDC to mask off bytes during byte transfers. For write cycles, the data masks have a latency of zero cycles, permitting data writes when the corresponding SDQMx pin is sampled low and blocking data writes when the SDQMx pin is sampled high on a byte-by-byte basis.

Executing a Parallel Refresh Command

The SDC includes a separate address pin (SA10) to enable the execution of Auto-Refresh commands in parallel with any asynchronous memory

access. This separate pin allows the SDC to issue a Precharge command to the SDRAM before it issues an Auto-Refresh command. In addition, the SA10 pin allows the SDC to enter and exit Self-Refresh mode in parallel with any asynchronous memory access.

The SA10 pin should be directly connected to the A10 pin of the SDRAM. During the Precharge command, SA10 is used to indicate that a Precharge All should be done. During a Bank Activate command, SA10 outputs the internal row address bit, which should be multiplexed to the A10 SDRAM input. During Read and Write commands, SA10 is used to disable the auto-precharge function of SDRAMs.

Since the ADSP-BF561 uses byte addressing, the SA10 pin replaces the ADDR11 pin for the 16-bit interface or ADDR12 pin for the 32-bit interface.

SDRAM Memory Global Control Register (EBIU_SDGCTL)

The SDRAM Memory Global Control register includes all programmable parameters associated with the SDRAM access timing and configuration. Figure 16-11 shows the EBIU_SDGCTL register bit definitions.

The SCTLE bit is used to enable or disable the SDC. If SCTLE is disabled, any access to SDRAM address space generates an internal bus error, and the access does not occur externally. For more information, see “Error Detection” on page 16-8. When SCTLE is disabled, all SDC control pins are in their inactive states and the SDRAM clock is not running. The SCTLE bit must be enabled for SDC operation and is enabled by default at reset.

To support higher clock load requirements, two SDRAM clock pins (SCLK0 and SCLK1) are provided. The clock timing of these two pins is identical, and the clock load on SCLK0 and SCLK1 should be balanced. SCLK0 runs whenever the SDC is enabled (SCTLE = 1). SCLK1 runs

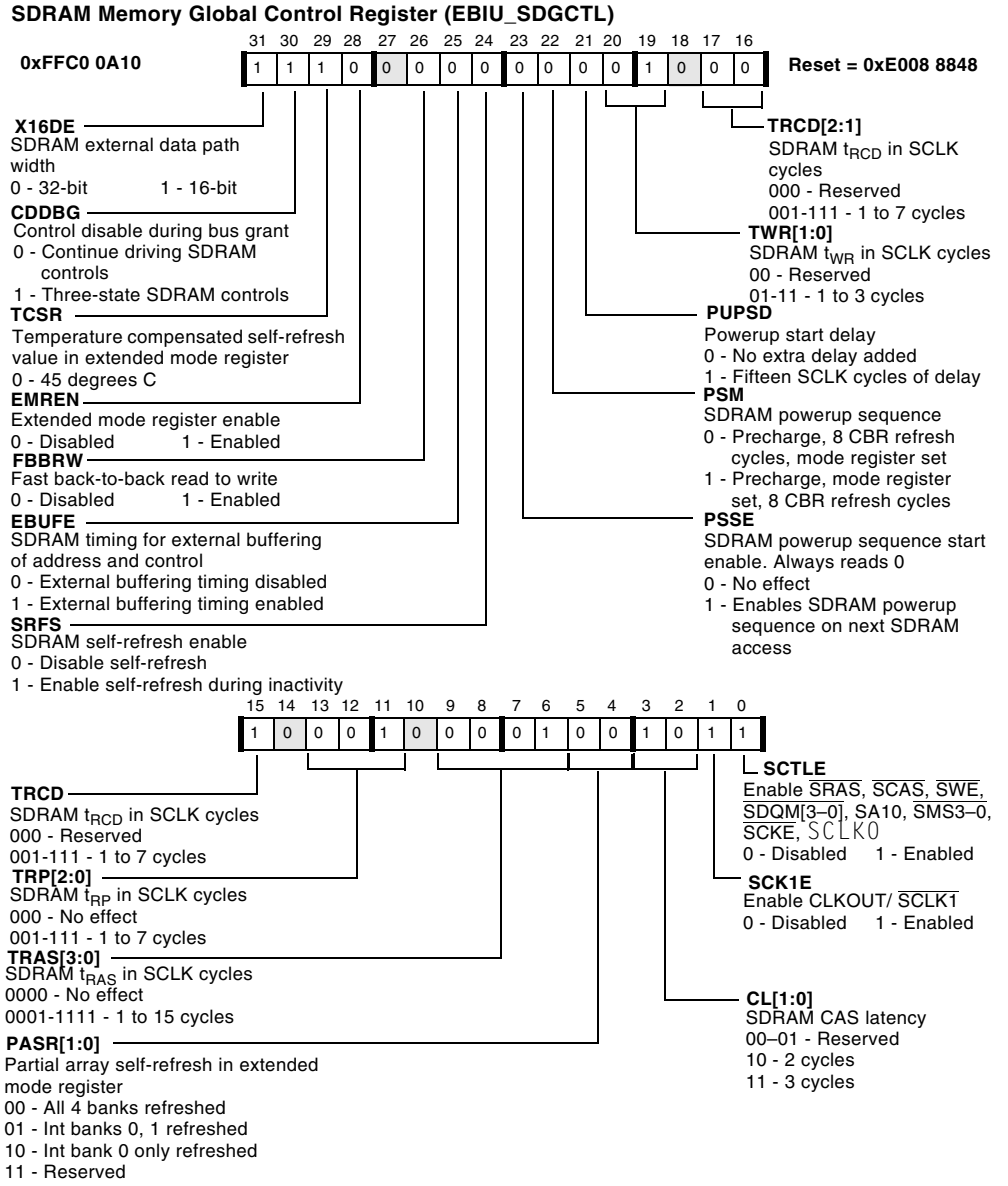

Figure 16-11. SDRAM Memory Global Control Register

whenever SCK1E is enabled. Both the SCTLE and SCK1E bits are enabled during reset, so SCLK0 and SCLK1 are running after reset deasserts. If the SDC will not be used, the SCTLE and SCK1E bits can be disabled to stop the clocks and reduce power dissipation. Even though the SDC is enabled (SCTLE = 1) at reset, the power-up sequence (PSSE = 1) must be executed before reading or writing to SDRAM address space.

The CAS Latency (CL), SDRAM t_{RAS} Timing (TRAS), SDRAM t_{RP} Timing (TRP), SDRAM t_{RCD} Timing (TRCD), and SDRAM t_{WR} Timing (TWR) bits should be programmed based on the system clock frequency and the timing specifications of the SDRAM used. Note all SDRAM banks use the same timing.

Note these timing parameters should not be changed while the SDC is active. Therefore, since the SDC is enabled by default upon reset, the SDC must be disabled after reset to allow for modifications to the default EBIU_SDGCTL bit values.

The PSM and PSSE bits work together to specify and trigger an SDRAM powerup (initialization) sequence. If the PSM bit is set to 1, the SDC does a Precharge All command, followed by a Load Mode Register command, and then does eight Auto-Refresh cycles. If the PSM bit is cleared, the SDC does a Precharge All command, followed by eight Auto-Refresh cycles, and then a Load Mode Register command. Two events must occur before the SDC does the SDRAM powerup sequence:

- The PSSE bit must be set to 1 to enable the SDRAM powerup sequence.
- A read or write access must be done to enabled SDRAM address space in order to have the external bus granted to the SDC so that the SDRAM powerup sequence may occur.

The SDRAM powerup sequence occurs and is followed immediately by the read or write transfer to SDRAM that was used to trigger the SDRAM

powerup sequence. Note there is a latency for this first access to SDRAM because the SDRAM powerup sequence takes many cycles to complete.

Before executing the SDC powerup sequence, ensure that the SDRAM receives stable power and is clocked for the proper amount of time, as specified by the SDRAM specification.

The Powerup Start Delay bit (PUPSD) optionally delays the powerup start sequence for 15 SCLK cycles. This is useful for multiprocessor systems sharing an external SDRAM. If the bus has been previously granted to the other processor before powerup and Self-Refresh mode is used when switching bus ownership, then the PUPSD bit can be used to guarantee a sufficient period of inactivity from self-refresh to the first Precharge command in the powerup sequence in order to meet the exit self-refresh time (t_{XSR}) of the SDRAM.

When the SRFS bit is set to 1, Self-Refresh mode is triggered. Once the SDC completes any active transfers, the SDC executes the sequence of commands to put the SDRAM into Self-Refresh mode. The next access to an enabled SDRAM bank causes the SDC to execute the commands to exit the SDRAM from Self-Refresh and execute the access. See "Entering and Exiting Self-Refresh Mode (SRFS)" on page 16-40 for more information about the SRFS bit.

The EBUFE bit is used to enable or disable external buffer timing. When buffered SDRAM modules or discrete register-buffers are used to drive the SDRAM control inputs, EBUFE should be set to 1. Using this setting adds a cycle of data buffering to read and write accesses. See "Setting SDRAM Buffering Timing Option (EBUFE)" on page 16-41 for more information about the EBUFE bit.

The FBBRW bit enables an SDRAM read followed by write to occur on consecutive cycles. In many systems, this is not possible because the turn-off time of the SDRAM data pins is too long, leading to bus contention with the succeeding write from the processor. When this bit is 0, a clock cycle is inserted between read accesses followed immediately by write accesses.

The EMREN bit enables programming of the Extended Mode register during startup. The Extended Mode register is used to control SDRAM power consumption in certain mobile low power SDRAMs. If the EMREN bit is enabled, then the TCSR and PASR1–0 bits control the value written to the Extended Mode register. The PASR bits determine how many SDRAM internal banks are refreshed during Self-Refresh. The TCSR bit signals to the SDRAM the worst case temperature range for the system, and thus how often the SDRAM internal banks need to be refreshed during Self-Refresh.

The CDDBG bit is used to enable or disable the SDRAM control signals when the external memory interface is granted to an external controller. If this bit is set to a 1, then the control signals are three-stated when bus grant is active. Otherwise, these signals continue to be driven during grant. If the bit is set and the external bus is granted, all SDRAM internal banks are assumed to have been changed by the external controller, so a precharge is required on each bank prior to use after control of the external bus is re-established. The control signals affected by this pin are $\overline{\text{SRAS}}$, $\overline{\text{SCAS}}$, $\overline{\text{SWE}}$, $\overline{\text{SMS3–0}}$, SA10, SCKE, and CLKOUT.

The X16DE bit is used to select whether the SDRAM interface is 32 bits wide or 16 bits wide. If X16DE is 0, DATA[31:0] should be connected to the SDRAM. If X16DE is 1, DATA[15:0] should be connected to the SDRAM. Note that all SDRAM banks must be either all 32 bits wide or all 16 bits wide.

Note all reserved bits in this register must always be written with 0s.

Setting SDRAM Clock Enables (SCTLE and SCK1E)

To meet higher clock load requirements for systems with multiple SDRAM devices, the ADSP-BF561 processor provides two SDRAM clock control pins, SCLK0 and SCLK1. These pins eliminate the need for off-chip clock buffers for most system memory configurations. The SCTLE and SCK1E bits in the SDRAM Memory Global Control register (EBIU_SDGCTL) provide control for the SDRAM clock control pins.

The SCTLE bit disables all of the SDRAM control pins: SDQM3–0, $\overline{\text{SCAS}}$, $\overline{\text{SRAS}}$, $\overline{\text{SWE}}$, SCKE, and the SCLK0.

- SCTLE = 0

  Disable all SDRAM control pins (control pins deasserted, SCLK0 low)

- SCTLE = 1

  Enable all SDRAM control pins (SCLK0 toggles)

The SCK1E bit disables the SCLK1 pin independently:

- SCK1E = 0

  Disable SCLK1 (SCLK1 low)

- SCK1E = 1

  Enable SCLK1 (SCLK1 toggles)

Note the SCLK1 function is also shared with the Asynchronous Memory Controllers' (AMC) CLKOUT function. Even if SCTLE and SCK1E are disabled, SCLK1 can be enabled independently by the CLKOUT enable in the AMC (AMCKEN in the EBIU_AMGCTL register).

If the system does not use SDRAM, both SCTLE and SCK1E should be 0.

If the system uses SDRAM, but the clock load is minimal, SCTLE should be 1 and SCK1E should be 0. This setting enables the SCLK0 pin and all related SDRAM control pins, but disables (holds low) the second clock pin SCLK1.

If the system uses SDRAM and the clock loading is high, both SCTLE and SCK1E should be set to 1. This setting enables SCLK0, SCLK1, and all SDRAM control pins. In this configuration, SCLK0 and SCLK1 should each share half of the clock load.

If an access occurs to the SDRAM address space while SCTLE is 0, the access generates an internal bus error and the access does not occur externally. For more information, see “Error Detection” on page 16-8. With careful software control, the SCTLE bit can be used in conjunction with Self-Refresh mode to further lower power consumption. However, SCTLE must remain enabled at all times when the SDC is needed to generate Auto-Refresh commands to SDRAM.

Entering and Exiting Self-Refresh Mode (SRFS)

The SDC supports SDRAM Self-Refresh mode. In Self-Refresh mode, the SDRAM performs refresh operations internally—without external control—reducing the SDRAM’s power consumption.

The SRFS bit in EBIU_SDGCTL enables the start of Self-Refresh mode:

- SRFS = 0

  Disable Self-Refresh mode

- SRFS = 1

  Enable Self-Refresh mode

When SRFS is set to 1, once the SDC enters an idle state it issues a Precharge command if necessary, and then issues a Self-Refresh command. If an internal access is pending, the SDC delays issuing the Self-Refresh command until it completes the pending SDRAM access and any subsequent pending access requests. Refer to “SDC Commands” on page 16-61 for more information.

Once the SDRAM device enters into Self-Refresh mode, the SDRAM controller asserts the SDSRA bit in the SDRAM Control Status register (EBIU_SDSTAT).

Note once the SRFS bit is set to 1, the SDC enters Self-Refresh mode when it finishes pending accesses. There is no way to cancel the entry into Self-Refresh mode.

The SDRAM device exits Self-Refresh mode only when the SDC receives a core or DMA access request. In conjunction with the SRFS bit, there are two ways to exit the self-refresh mode:

If the SRFS bit is set before the request, the SDC exits self-refresh and remains in auto-refresh mode.

If the SRFS bit is cleared before the request, the SDC exits self-refresh only for a single request and returns back to self-refresh mode until a new request arrives.

When the device comes out of reset, it will exit self-refresh mode and the SCKE pin will be driven high. If it is necessary to remain in self-refresh mode after reset, the SCKE may be programmed to stay low. Read the discussion on page 8-25 for details on using the SCKELOW bit to program the SCKE pin.

Setting SDRAM Buffering Timing Option (EBUFE)

To meet overall system timing requirements, systems that employ several SDRAM devices connected in parallel may require buffering between the processor and multiple SDRAM devices. This buffering generally consists of a register and driver.

To meet such timing requirements and to allow intermediary registration, the SDC supports pipelining of SDRAM address and control signals.

The EBUFE bit in the EBIU_SDGCTL register enables this mode:

- EBUFE = 0

  Disable external buffering timing

- EBUFE = 1

  Enable external buffering timing

When EBUFE = 1, the SDRAM controller delays the data in write accesses by one cycle, enabling external buffer registers to latch the address and controls. In read accesses, the SDRAM controller samples data one cycle later to account for the one-cycle delay added by the external buffer registers. When external buffering timing is enabled, the latency of all accesses is increased by one cycle.

Selecting the CAS Latency Value (CL)

The CAS latency value defines the delay, in number of clock cycles, between the time the SDRAM detects the Read command and the time it provides the data at its output pins.

CAS latency does not apply to write cycles.

The CL bits in the SDRAM Memory Global Control register (EBIU_SDGCTL) select the CAS latency value:

- CL = 00

  Reserved

- CL = 01

  Reserved

- `CL = 10`

  2 clock cycles

- `CL = 11`

  3 clock cycles

Generally, the frequency of operation determines the value of the CAS latency. For specific information about setting this value, consult the SDRAM device documentation.

Selecting the RAS to CAS Delay (TRCD)

The t_{RCD} value (RAS to CAS delay) defines the delay for the first read or write command after a row activate command, in number of clock cycles. The t_{RCD} parameter allows the processor to adapt to the timing requirements of the system's SDRAM devices. The `TRCD` bits in the SDRAM Memory Global Control register (`EBIU_SDGCTL`) select the t_{RCD} value. Any value between 1 and 7 clock cycles may be selected. For example:

`TRCD` = reserved	no effect
`TRCD` = 001	1 clock cycle
`TRCD` = 010	2 clock cycles
`TRCD` = 111	7 clock cycles

Selecting the Bank Activate Command Delay (TRAS)

The t_{RAS} value (Bank Activate command delay) defines the required delay, in number of clock cycles, between the time the SDC issues a Bank Activate command and the time it issues a Precharge command. The SDRAM must also remain in Self-Refresh mode for at least the time period specified by t_{RAS}. The t_{RP} and t_{RAS} values define the t_{RFC}, t_{RC}, and t_{XSR} values. More information can be found on page 16-30.

The t_{RAS} parameter allows the processor to adapt to the timing requirements of the system's SDRAM devices.

The TRAS bits in the SDRAM Memory Global Control register (EBIU_SDGCTL) select the t_{RAS} value. Any value between 1 and 15 clock cycles can be selected. For example:

- TRAS = 0000

  No effect

- TRAS = 0001

  1 clock cycle

- TRAS = 0010

  2 clock cycles

- TRAS = 1111

  15 clock cycles

For specific information on setting this value, consult the SDRAM device documentation.

Selecting the Precharge Delay (TRP)

The t_{RP} value (Precharge delay) defines the required delay, in number of clock cycles, between the time the SDC issues a Precharge command and the time it issues a Bank Activate command. The t_{RP} also specifies the time required between Precharge and Auto-Refresh, and between Precharge and Self-Refresh. The t_{RP} and t_{RAS} values define the t_{RFC}, t_{RC}, and t_{XSR} values.

This parameter enables the application to accommodate the SDRAM's timing requirements.

The TRP bits in the SDRAM Memory Global Control register (EBIU_SDGCTL) select the t_{RP} value. Any value between 1 and 7 clock cycles may be selected. For example:

- TRP = 000

  No effect

- TRP = 001

  1 clock cycle

- TRP = 010

  2 clock cycles

- TRP = 111

  7 clock cycles

Selecting the Write to Precharge Delay (TWR)

The t_{WR} value defines the required delay, in number of clock cycles, between the time the SDC issues a Write command (drives write data) and a Precharge command.

This parameter enables the application to accommodate the SDRAM's timing requirements.

The TWR bits in the SDRAM Memory Global Control register (EBIU_SDGCTL) select the t_{WR} value. Any value between 1 and 3 clock cycles may be selected. For example:

- TWR = 00

  Reserved

- `TWR = 01`

  1 clock cycle

- `TWR = 10`

  2 clock cycles

- `TWR = 11`

  3 clock cycles

SDRAM Memory Bank Control Register (EBIU_SDBCTL)

The SDRAM Memory Bank Control register (Figure 16-12) includes individual external bank-specific programmable parameters. It allows software to control some parameters of the SDRAM. The external bank can be configured for a different size of SDRAM. It uses the access timing parameters defined in the SDRAM Memory Global Control register (`EBIU_SDGCTL`). The `EBIU_SDBCTL` register should be programmed before powerup and should be changed only when the SDC is idle.

Each bank is independently programmable and is contiguous with the adjacent bank regardless of the sizes of the different banks or their placements.

The `EBIU_SDBCTL` register stores the configuration information for the SDRAM bank interface. The EBIU supports 64 Mbit, 128 Mbit, 256 Mbit, and 512 Mbit SDRAM devices with x4, x8, x16 configurations. See "SDRAM External Bank Address Decode" on page 16-52 for more information on bank starting address decodes.

The SDC determines the internal SDRAM page size from the `EBCAW` parameters. Page sizes of 512 byte, 1K byte, 2K byte, and 4K byte are supported for 16-bit SDRAM banks, while 1, 2, 4, or 8 KB are supported for 32-bit SDRAM. Table 16-7 and Table 16-8 show the page size and

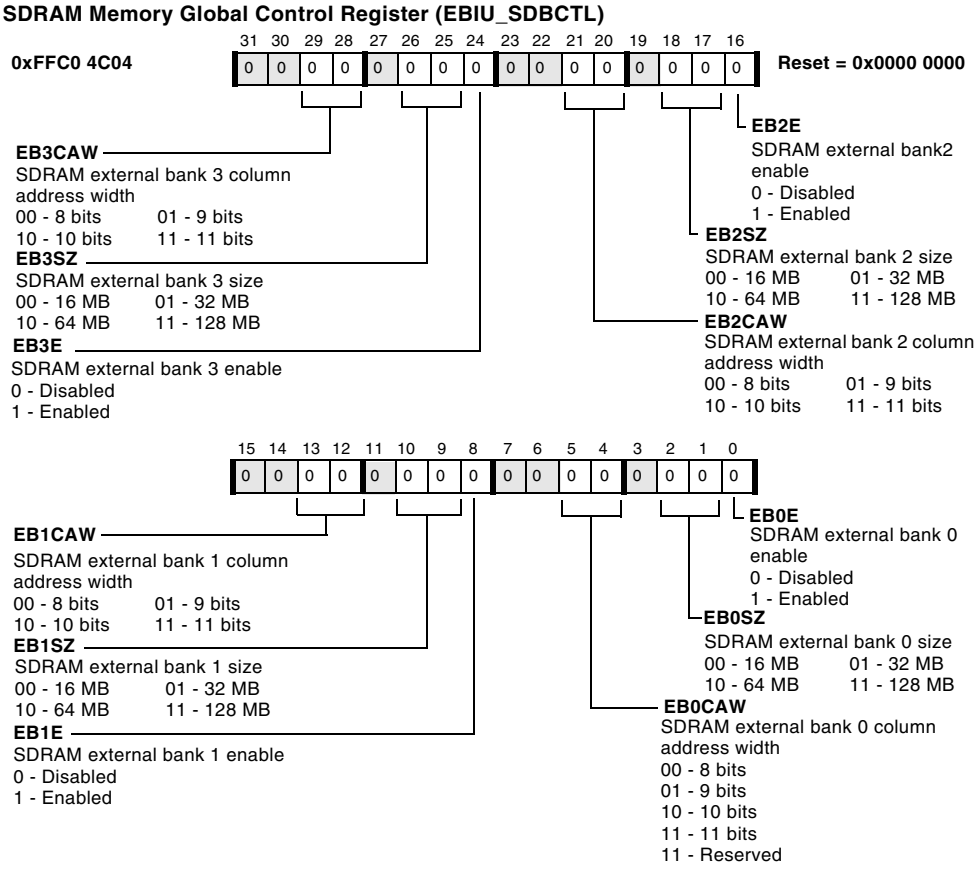

Figure 16-12. SDRAM Memory Bank Control (EBIU_SDBCTL)

breakdown of the internal address (IA31–0, as seen from the core or DMA) into the row, bank, column, and byte address. The column address and the byte address together make up the address inside the page.

The EBE bits in the EBIU_SDBCTL register are used to enable or disable the external SDRAM bank. If the SDRAM is disabled, any access to the SDRAM address space generates an internal bus error, and the access does not occur externally. For more information, see “Error Detection” on page 16-8.

Table 16-7. Internal Address Mapping (16-Bit Bank Width Example)

Bank Width (bits)	Bank Size (M byte)	Col. Addr. Width (CAW)	Page Size (K byte)	Bank Address	Row Address	Page	
						Column Address	Byte Address
16	128	11	4	IA26—25	IA24—12	A11—1	IA0
16	128	10	2	IA26—25	IA24—11	IA10—1	IA0
16	128	9	1	1A26—25	IA24—10	IA9—1	IA0
16	128	8	.5	IA26—25	IA24—9	IA8—1	IA0
16	64	11	4	IA25—24	IA23—12	IA11—1	IA0
16	64	10	2	IA25—24	IA23—11	IA10—1	IA0
16	64	9	1	IA25—24	IA23—10	IA9—1	IA0
16	64	8	.5	IA25—24	IA23—9	IA8—1	IA0
16	32	11	4	IA24—23	IA22—12	IA11—1	IA0
16	32	10	2	IA24—23	IA22—11	IA10—1	IA0
16	32	9	1	IA24—23	IA22—10	IA9—1	IA0
16	32	8	.5	IA24—23	IA22—9	IA8—1	IA0
16	16	11	4	IA23—22	IA21—12	IA11—1	IA0
16	16	10	2	IA23—22	IA21—11	IA10—1	IA0
16	16	9	1	IA23—22	IA21—10	IA9—1	IA0
16	16	8	.5	IA23—22	IA21—9	IA8—1	IA0

Table 16-8. Internal Address Mapping (32-Bit Bank Width Example)

Bank Width (bits)	Bank Size (Mbyte)	Col. Addr. Width (CAW)	Page Size (kbyte)	Bank Address	Row Address	Page	
						Column Address	Byte Address
32	128	11	8	IA26—25	IA24—13	A12—2	IA1—0
32	128	10	4	IA26—25	IA24—12	IA11—2	IA1—0
32	128	9	2	1A26—25	IA24—11	IA10—2	IA1—0
32	128	8	1	IA26—25	IA24—10	IA9—2	IA1—0
32	64	11	8	IA25—24	IA23—13	A12—2	IA1—0
32	64	10	4	IA25—24	IA23—12	IA11—2	IA1—0
32	64	9	2	IA25—24	IA23—11	IA10—2	IA1—0
32	64	8	1	IA25—24	IA23—10	IA9—2	IA1—0
32	32	11	8	IA24—23	IA22—13	A12—2	IA1—0
32	32	10	4	IA24—23	IA22—12	IA11—2	IA1—0
32	32	9	2	IA24—23	IA22—11	IA10—2	IA1—0
32	32	8	1	IA24—23	IA22—10	IA9—2	IA1—0
32	16	11	8	IA23—22	IA21—13	A12—2	IA1—0
32	16	10	4	IA23—22	IA21—12	IA11—2	IA1—0
32	16	9	2	IA23—22	IA21—11	IA10—2	IA1—0
32	16	8	1	IA23—22	IA21—10	IA9—2	IA1—0

SDRAM Control Status Register (EBIU_SDSTAT)

The SDRAM Control Status register (Figure 16-13) provides information on the state of the SDC. This information can be used to determine when it is safe to alter SDC control parameters or it can be used as a debug aid. The SDEASE bit of this register is sticky. Once it has been set, software must explicitly write a 1 to the bit to clear it. Writes have no effect on the

other status bits, which are updated by the SDC only. This SDC MMR is 16 bits wide.

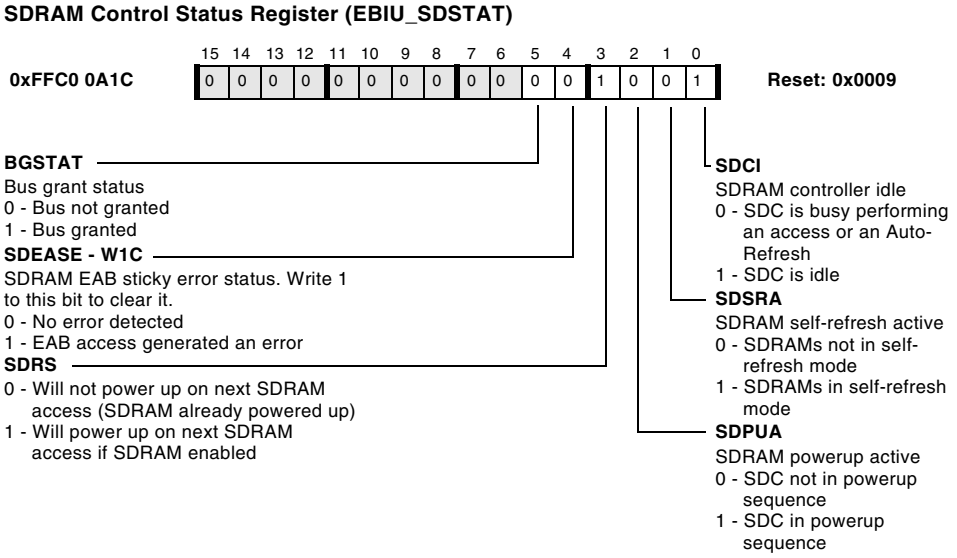

Figure 16-13. SDRAM Control Status (EBIU_SDSTAT)

SDRAM Refresh Rate Control Register (EBIU_SDRRC)

The SDRAM Refresh Rate Control register (Figure 16-14) provides a flexible mechanism for specifying the Auto-Refresh timing. Since the clock supplied to the SDRAM can vary, the SDC provides a programmable refresh counter, which has a period based on the value programmed into the RDIV field of this register. This counter coordinates the supplied clock rate with the SDRAM device's required refresh rate.

The desired delay (in number of SDRAM clock cycles) between consecutive refresh counter time-outs must be written to the RDIV field. A refresh counter time-out triggers an Auto-Refresh command to all external SDRAM devices. Write the RDIV value to the EBIU_SDRRC register before

the SDRAM powerup sequence is triggered. Change this value only when the SDC is idle.

SDRAM Refresh Rate Control Register (EBIU_SDRRC)

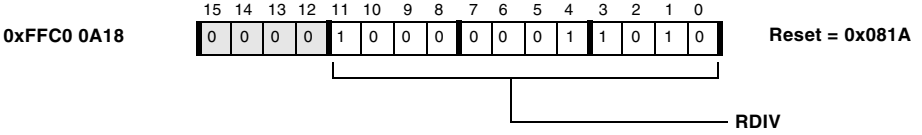

Figure 16-14. SDRAM Refresh Rate Control (EBIU_SDRRC)

To calculate the value that should be written to the EBIU_SDRRC register, use the following equation:

$RDIV = ((f_{SCLK} \times t_{REF}) / NRA) - (t_{RAS} + t_{RP})$

Where:

- f_{SCLK} = SDRAM clock frequency (system clock frequency)
- t_{REF} = SDRAM refresh period
- NRA = Number of row addresses in SDRAM (refresh cycles to refresh whole SDRAM)
- t_{RAS} = Active to Precharge time (TRAS in the SDRAM Memory Global Control register) in number of clock cycles
- t_{RP} = RAS to Precharge time (TRP in the SDRAM Memory Global Control register) in number of clock cycles

This equation calculates the number of clock cycles between required refreshes and subtracts the required delay between Bank Activate commands to the same internal bank ($t_{RC} = t_{RAS} + t_{RP}$). The t_{RC} value is subtracted, so that in the case where a refresh time-out occurs while an SDRAM cycle is active, the SDRAM refresh rate specification is

guaranteed to be met. The result from the equation should always be rounded down to an integer.

Below is an example of the calculation of RDIV for a typical SDRAM in a system with a 133 MHz clock:

- f_{SCLK} = 133 MHz
- t_{REF} = 64 ms
- NRA = 4096 row addresses
- t_{RAS} = 2
- t_{RP} = 2

The equation for RDIV yields:

- RDIV = $((133 \times 10^6 \times 64 \times 10^{-3}) / 4096) - (2 + 2) = 2074$ clock cycles

This means RDIV is 0x81A (hex) and the SDRAM Refresh Rate Control register should be written with 0x081A.

Note RDIV must be programmed to a nonzero value if the SDRAM controller is enabled. When RDIV = 0, operation of the SDRAM controller is not supported and can produce undesirable behavior. Values for RDIV can range from 0x001 to 0xFFF.

SDRAM External Bank Address Decode

The memory address space for each SDRAM bank is contiguous. Therefore, the starting addresses for each of the SDRAM banks and for the beginning of the unpopulated/reserved space is determined by adding up the sizes of each bank that precedes that bank. This can be done by simply summing the Bank Size Encodings listed in Table 16-9 for each bank as shown in Table 16-10. If every SDRAM bank is 128 MB, the SDRAM address space is fully populated. If any bank is less than 128 MB, the

address space after bank3 is referred to as “unpopulated/reserved” space. Any access to unpopulated/reserved space generates an internal bus error and the access does not occur externally.

Table 16-9. Bank Size Encodings

EBxSZ	Bank Size (Mbyte)	bank_x_size[4:0]
00	16	00001
01	32	00010
10	64	00100
11	128	01000

Table 16-10. Start Address Calculation

SDRAM Bank	bank_x_start_addr[28:24]
0	00000
1	00000 + bank_0_size[4:0]
2	00000 + bank_0_size[4:0] + bank_1_size[4:0]
3	00000 + bank_0_size[4:0] + bank_1_size[4:0] + bank_2_size[4:0]
Unpopulated/ Reserved	00000 + bank_0_size[4:0] + bank_1_size[4:0] + bank_2_size[4:0] + bank_3_size[4:0]

Note:
1. bank_x_start_addr31—29 = 000 and bank_x_start_addr23—0 = all 0’s
2. If there is a carry out on the reserved space start address, then the SDRAM address space is fully populated and therefore, no bus error from unpopulated/reserved space is possible.

For example, if the SDRAM memory configuration is as follows:

- Bank0 and bank1 are connected to a 64 MByte SDRAM DIMM (EB0E=1, EB0SZ=01, EB1E=1, EB1SZ=01).
- Bank2 is unpopulated, but is set to be 16 MB (EB2E=0, EB2SZ=00).
- Bank3 is connected to a 64 MByte discrete SDRAM memory array (EB3E=0, EB3SZ=10).

```
bank_0_start_addr28—24 =    00000

bank_1_start_addr28—24 =    00000
                            00010      bank0
                            -----
                            00010

bank_2_start_addr28—24 =    00000
                            00010      bank0
                            00010      bank1
                            -----
                            00100

bank_3_start_addr28—24 =    00000
                            00010      bank0
                            00010      bank1
                            00001      bank2
                            -----
                            00101

unpopulated_start_addr28—24 = 00000
                            00010      bank0
                            00010      bank1
                            00001      bank2
                            00100      bank3
                            -----
                            00101
```

Figure 16-15. Bank Start Address Calculation

The start addresses for the SDRAM banks and reserved spaces are calculated as:

For all banks, bank_x_start_addr[31:29] = 000 and bank_x_start_addr[32:0] = all zeros. Therefore, the bank size encodings are as shown in Table 16-11.

Note even though bank2 is disabled (EB2E=0), the bank size (EB2SZ=00) is used to determine the start addresses of successive banks and unpopulated/reserved space.

Table 16-11. Bank Size Encodings

SDRAM Bank	Bank Size (MB)	Start Address (hex)
0	32	0x0000 0000
1	32	0x0200 0000
2	16	0x0400 0000
3	64	0x0500 0000
Reserved	368	0x0900 0000

SDRAM Address Mapping

To access SDRAM, the SDC multiplexes the internal 32-bit non-multiplexed address into a row address, a column address, a bank address, and the byte data masks for the SDRAM device. See Figure 16-16. The lowest bit or bits are mapped to byte data masks, the next bits are mapped into the column address, the next bits are mapped into the row address, and the final two bits are mapped into the bank address. This mapping is based on the EBxSZ and EBxCAW parameters programmed into the SDRAM Memory Bank Control register, and the X16DE bit in the EBIU_SDGTL register.

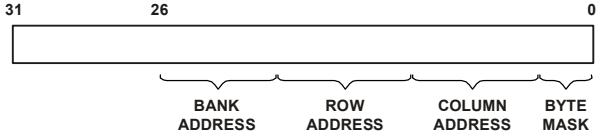

Figure 16-16. Multiplexed SDRAM Addressing Scheme

There is no separate ADDR1 pin on the ADSP-BF561 processor. The ADDR1 pin is muxed onto the pin $\overline{\text{ABE[3]}}$/$\overline{\text{SDQM[3]}}$/ADDR1 for 16-bit ASYNC/SDRAM banks. This is further described in "Data Mask (SDQM3—0) Encodings" on page 16-57.

16-Bit Wide SDRAM Address Muxing

Table 16-12 shows the connection of the address pins with the SDRAM device pins.

Table 16-12. SDRAM Address Connections for 16-Bit and 32-Bit Banks

SDRAM Address Pin	External Address Pin for 16-Bit Mode	External Address Pin for 32-Bit Mode
BA1	ADDR19	ADDR19
BA0	ADDR18	ADDR18
A15	ADDR16	ADDR17
A14	ADDR15	ADDR16
A13	ADDR14	ADDR15
A12	ADDR13	ADDR14
A11	ADDR12	ADDR13
A10	SA10	SA10
A9	ADDR10	ADDR11
A8	ADDR9	ADDR10
A7	ADDR8	ADDR9
A6	ADDR7	ADDR8
A5	ADDR6	ADDR7
A4	ADDR5	ADDR6
A3	ADDR4	ADDR5
A2	ADDR3	ADDR4
A1	ADDR2	ADDR3
A0	SDQM3 (ADDR1)	ADDR2

Data Mask (SDQM3—0) Encodings

During write transfers to SDRAM, the SDQM3–0 pins are used to mask writes to bytes that are not accessed. Table 16-13 shows the SDQM3–0 encodings for 32-bit wide SDRAM banks based on the internal transfer address bits IA1–0 and the transfer size.

During read transfers to SDRAM banks, reads are always done of all bytes in the bank regardless of the transfer size. This means for 32-bit SDRAM banks, SDQM3–0 are all zeros and for 16-bit SDRAM banks, SDQM1–0 are zeros, SDQM2 is 1, and SDQM3 outputs IA1.

The only time that the SDQM3–0 pins are high is when bytes are masked during write transfers to the SDRAM banks. At all other times, the SDQM3–0 pins are held low (with the exception of SDQM3 when the SDRAM banks are 16-bit wide). This means for 32-bit SDRAM banks, SDQM3–0 are all zeros, and for 16-bit SDRAM banks SDQM1–0 are zeros, SDQM2 is 1, and SDQM3 outputs IA1.

The pin configurations in Table 16-13 are internally set up and not programmable.

Table 16-13. SDQM3—0 Encodings During Writes for 32-Bit SDRAM Banks

Internal Address IA[1:0]	Internal Transfer Size		
	byte	2 bytes	4 bytes
00	SDQM3=1 SDQM2=1 SDQM1=1 SDQM0=0	SDQM3=1 SDQM2=1 SDQM1=0 SDQM0=0	SDQM3=0 SDQM2=0 SDQM1=0 SDQM0=0
01	SDQM3=1 SDQM2=1 SDQM1=0 SDQM0=1		

Table 16-13. SDQM3—0 Encodings During Writes for 32-Bit SDRAM Banks (Cont'd)

Internal Address IA[1:0]	Internal Transfer Size		
	byte	2 bytes	4 bytes
10	SDQM3=1 SDQM2=0 SDQM1=1 SDQM0=1	SDQM3=0 SDQM2=0 SDQM1=1 SDQM0=1	
11	SDQM3=0 SDQM2=1 SDQM1=1 SDQM0=1		

Table 16-14 shows the SDQM1—0 encodings during writes for 16-bit SDRAM banks.

Table 16-14. SDQM1—0 Encodings During Writes for 16-Bit SDRAM Banks

Internal Address IA[0]	Internal Transfer Size		
	byte	2 bytes	4 bytes
0	SDQM1=1 SDQM0=0	SDQM1=0 SDQM0=0	SDQM1=0 SDQM0=0
1	SDQM1=0 SDQM0=1		
Note: SDQM[3] **outputs** IA[1].			

SDC Operation

The SDC uses a burst length = 1 for read and write operations. Whenever a page miss occurs, the SDC executes a Precharge command followed by a Bank Activate command before executing the Read or Write command. If there is a page hit, the Read or Write command can be given immediately without requiring the Precharge command.

For SDRAM Read commands, there is a latency from the start of the Read command to the availability of data from the SDRAM, equal to the CAS latency. This latency is always present for the first read and for any single read transfer. Subsequent reads in a burst do not have latency.

A programmable refresh counter is provided. It can be programmed to generate background Auto-Refresh cycles at the required refresh rate based on the clock frequency used. The refresh counter period is specified with the RDIV field in the SDRAM Refresh Rate Control register.

To allow Auto-Refresh commands to execute in parallel with any AMC access, a separate A10 pin (SA10) is provided. All the SDRAM internal banks are precharged before issuing an Auto-Refresh command.

The internal 32-bit non-multiplexed address is multiplexed into a row address, a column address, a bank select address, and data masks. Bit0 for 16-bit wide SDRAMs is used to generate the data masks. The next lowest bits are mapped into the column address, next bits are mapped into the row address, and the final two bits are mapped into the internal bank address. This mapping is based on the EBCAW and EBSZ values programmed into the SDRAM Memory Bank Control register.

SDC Configuration

After a processor's hardware or software reset, the SDC clocks are enabled; however, the SDC must be configured and initialized. Before programming the SDC and executing the powerup sequence, ensure the clock to the SDRAM is enabled after the power has stabilized for the proper amount of time (as specified by the SDRAM). In order to set up the SDC and start the SDRAM powerup sequence for the SDRAMs, the SDRAM Refresh Rate Control register (EBIU_SDRRC), the SDRAM Memory Bank Control register (EBIU_SDBCTL), and SDRAM Memory Global Control register (EBIU_SDGCTL) must be written, and a transfer must be started to SDRAM address space. The SDRS bit of the SDRAM Control Status

register can be checked to determine the current state of the SDC. If this bit is set, the SDRAM powerup sequence has not been initiated.

The RDIV field of the EBIU_SDRRC register should be written to set the SDRAM refresh rate.

The EBIU_SDBCTL register should be written to describe the sizes and SDRAM memory configurations used (EBxSZ and EBxCAW) and to enable the external bank (EBEx). Note until the SDRAM powerup sequence has been started, any access to SDRAM address space, regardless of the state of the EBE bit, generates an internal bus error, and the access does not occur externally. For more information, see "Error Detection" on page 16-8. After the SDRAM powerup sequence has completed, if the external bank is disabled, any transfer to it results in a hardware error interrupt, and the SDRAM transfer does not occur.

The EBIU_SDGCTL register is written:

- to set the SDRAM cycle timing options (CL, TRAS, TRP, TRCD, TWR, EBUFE)
- to enable the SDRAM clock (SCTLE, SCK1E)
- to select and enable the start of the SDRAM powerup sequence (PSM, PSSE)

Note if SCTLE is disabled, any access to SDRAM address space generates an internal bus error and the access does not occur externally. For more information, see "Error Detection" on page 16-8.

Once the PSSE bit in the EBIU_SDGCTL register is set to 1, and a transfer occurs to enabled SDRAM address space, the SDC initiates the SDRAM powerup sequence. The exact sequence is determined by the PSM bit in the EBIU_SDGCTL register. The transfer used to trigger the SDRAM powerup sequence can be either a read or a write. This transfer occurs when the SDRAM powerup sequence has completed. This initial transfer takes

many cycles to complete since the SDRAM powerup sequence must take place.

SDC Commands

This section provides a description of each of the commands that the SDC uses to manage the SDRAM interface. These commands are initiated automatically upon a memory read or memory write. A summary of the various commands used by the on-chip controller for the SDRAM interface is as follows.

- Single precharge: Precharge a single bank
- Precharge all: Precharge all banks
- Extended mode register: Initializes the mobile SDRAMs operation parameters during the power-up sequence
- Bank Activate: Activates a page in the required SDRAM internal bank
- Load Mode Register: Initializes the SDRAM operation parameters during the power-up sequence
- Read/Write
- Auto-Refresh: Causes the SDRAM to execute a CAS before RAS refresh
- Self-Refresh: Places the SDRAM in self-refresh mode, in which the SDRAM powers down and controls its refresh operations internally
- NOP/Command Inhibit: No operation

Table 16-15 shows the SDRAM pin state during SDC commands.

Table 16-15. Pin State During SDC Commands

Command	$\overline{\text{SMS3–0}}$	$\overline{\text{SCAS}}$	$\overline{\text{SRAS}}$	$\overline{\text{SWE}}$	SCKE	SA10
Precharge All	low	high	low	low	high	high
Bank Activate	low	high	low	high	high	
Load Mode Register	low	low	low	low	high	
Read	low	low	high	high	high	low
Write	low	low	high	low	high	low
Auto-Refresh	low	low	low	high	high	
Self-Refresh	low	low	low	high	low	
NOP	low	high	high	high	high	
Command Inhibit	high	high	high	high	high	
Single Precharge	low	high	low	low	high	low

Precharge Command

The Precharge All command is given to precharge all internal banks at the same time before executing an auto-refresh. For a page miss during reads or writes in a specific internal SDRAM bank, the SDC uses the Single Precharge command to that bank.

Bank Activate Command

The Bank Activate command is required if the next data access is in a different page of the same SDRAM bank. The SDC executes the Precharge command, followed by a Bank Activate command, to activate the page in the desired page in that bank.

Since an SDRAM contains four independent banks (A–D), the SDC can take advantage of the architecture and support bank interleaving.

Any first access to SDRAM bank A will force an activate command before a read or write command. However, if any new access falls into the address space of the banks (B, C, D), the SDC leaves bank A open and activates any of the banks (B, C, D). This scenario is repetitive until all four banks (A–D) are opened. This results in an effective page size of up to four pages because no latency causes switching between these open pages (compared to one page in only one bank at the time).

Any access to any closed page in any opened bank (A-D) forces a precharge command to that bank only. The periodic auto-refresh command is preceded by a precharge-all command, where all open SDRAM banks are closed before the refresh cycle is performed.

The benefit of Bank interleaving reduces precharge and activation cycles by mapping instructions or data among different internal SDRAM banks.

Although there can be as many as 16 internal banks available in four external banks, only four can be open at a time. Since any internal bank (A-D) can be open in only one external SDRAM bank, if internal bank A is open in external bank 3, the external banks 0, 1, and 2 cannot have their internal bank A open. Here are some examples:

- If there is only one external SDRAM bank, then all of its internal banks can be open at the same time.
- External SDRAM bank 2 may have all of its internal banks open. All internal banks in external SDRAM banks 0, 1, and 3 are closed.
- Internal bank A open in external bank 3. Internal bank B open in external bank 1. Internal bank C open in external bank 2. Internal bank D open in external bank 0. All other internal banks are closed.

Load Mode Register Command

The Load Mode Register command initializes SDRAM operation parameters. This command is a part of the SDRAM powerup sequence. The Load Mode Register command uses the address bus of the SDRAM as data input. The powerup sequence is initiated by writing 1 to the PSSE bit in the SDRAM Memory Global Control register (EBIU_SDGCTL) and then writing or reading from any enabled address within the SDRAM address space to trigger the powerup sequence. The exact order of the powerup sequence is determined by the PSM bit of the EBIU_SDGCTL register.

The Load Mode Register command initializes these parameters:

- Burst length = 1, bits 2–0, always 0
- Wrap type = sequential, bit 3, always 0
- Ltmode = latency mode (CAS latency), bits 6–4, programmable in the EBIU_SDGCTL register
- Bits 14–7, always 0

While executing the Load Mode Register command, the unused address pins are set to 0. During the two clock cycles following the Load Mode Register command, the SDC issues only NOP commands.

For low power mobile SDRAMs that include an Extended Mode register, this register is programmed during powerup sequence if the EMREN bit is set in the EBIU_SDGCTL register.

The Extended Mode register is initialized with these parameters:

- Partial Array Self-Refresh, bits 2–0, bit 2 always 0, bits 1–0 programmable in EBIU_SDGCTL
- Temperature Compensated Self-Refresh, bits 4–3, bit 3 always 1, bit 4 programmable in EBIU_SDGCTL
- Bits 12–5, always 0, and bit 13 always 1

Read/Write Command

A Read/Write command is executed if the next read/write access is in the present active page. During the Read command, the SDRAM latches the column address. The delay between Activate and Read commands is determined by the t_{RCD} parameter. Data is available from the SDRAM after the CAS latency has been met.

In the Write command, the SDRAM latches the column address. The write data is also valid in the same cycle. The delay between Activate and Write commands is determined by the t_{RCD} parameter.

The SDC does not use the auto-precharge function of SDRAMs, which is enabled by asserting SA10 high during a Read or Write command.

Auto-Refresh Command

The SDRAM internally increments the refresh address counter and causes a CAS before RAS (CBR) refresh to occur internally for that address when the Auto-Refresh command is given. The SDC generates an Auto-Refresh command after the SDC refresh counter times out. The RDIV value in the SDRAM Refresh Rate Control register must be set so that all addresses are refreshed within the t_{REF} period specified in the SDRAM timing specifications. The command is issued to the enabled SDRAM bank. Before executing the Auto-Refresh command, the SDC executes a Precharge All command to the external bank. The next Activate command is not given until the t_{RFC} specification ($t_{RFC} = t_{RAS} + t_{RP}$) is met.

Auto-Refresh commands are also issued by the SDC as part of the powerup sequence and also after exiting Self-Refresh mode.

Self-Refresh Command

The Self-Refresh command causes refresh operations to be performed internally by the SDRAM, without any external control. This means that the SDC does not generate any Auto-Refresh cycles while the SDRAM is

in Self-Refresh mode. Before executing the Self-Refresh command, all internal banks are precharged. Self-Refresh mode is enabled by writing a 1 to the SRFS bit of the SDRAM Memory Global Control register (EBIU_SDGCTL). After issuing the Self-Refresh command, the SDC drives SCKE low. This puts the SDRAM into a power down mode (SCKE = 0, $\overline{\text{SRAS}}/\overline{\text{SMS}}/\overline{\text{SCAS}}/\overline{\text{SWE}}$ = $\overline{1}$) Before exiting Self-Refresh mode, the SDC asserts SCKE. The SDRAM remains in Self-Refresh mode for at least t_{RAS} and until an internal access to SDRAM space occurs. When an internal access occurs causing the SDC to exit the SDRAM from Self-Refresh mode, the SDC waits to meet the t_{XSR} specification ($t_{XSR} = t_{RAS} + t_{RP}$) and then issues an Auto-Refresh command. After the Auto-Refresh command, the SDC waits for the t_{RFC} specification ($t_{RFC} = t_{RAS} + t_{RP}$) to be met before executing the Activate command for the transfer that caused the SDRAM to exit Self-Refresh mode. Therefore, the latency from when a transfer is received by the SDC while in Self-Refresh mode, until the Activate command occurs for that transfer, is $2 \times (t_{RAS} + t_{RP})$.

Note CLKOUT is not disabled by the SDC during Self-Refresh mode. However, software may disable the clock by clearing the SCTLE or SCK1E bits in EBIU_SDGCTL. The application software should ensure that all applicable clock timing specifications are met before the transfer to SDRAM address space which causes the controller to exit Self-Refresh mode. If a transfer occurs to SDRAM address space when the SCTLE bit is cleared, an internal bus error is generated, and the access does not occur externally, leaving the SDRAM in Self-Refresh mode. For more information, see “Error Detection” on page 16-8.

After issuing the self refresh command, the SDC drives SCKE low. This puts the SDRAMs into a power down mode (SCKE = 0, $\overline{\text{SRAS}}/\overline{\text{SMS}}/\overline{\text{SCAS}}/\overline{\text{SWE}}$ = 0). Before exiting the self refresh mode, the SDC asserts SCKE.

When in the Deep Sleep operating mode, clocking to the SDRAM is turned off. Before entering the Deep Sleep mode, software should ensure that SDRAM is entered in self-refresh mode.

No Operation/Command Inhibit Commands

The No Operation (NOP) command to the SDRAM has no effect on operations currently in progress. The Command Inhibit command is the same as a NOP command; however, the SDRAM is not chip-selected. When the SDC is actively accessing the SDRAM but needs to insert additional commands with no effect, the NOP command is given. When the SDC is not accessing the SDRAM, the Command Inhibit command is given.

SDRAM Timing Specifications

To support key timing requirements and powerup sequences for different SDRAM vendors, the SDC provides programmability for t_{RAS}, t_{RP}, t_{RCD}, t_{WR}, and the powerup sequence mode. (For more information, see "SDRAM Memory Global Control Register (EBIU_SDGCTL)" on page 16-34.) CAS latency should be programmed in the `EBIU_SDGCTL` register based on the frequency of operation. (Please refer to the SDRAM vendor's data sheet for more information.)

For other parameters, the SDC assumes:

- Bank Cycle Time: $t_{RC} = t_{RAS} + t_{RP}$
- Row (Bank A) to Row (Bank B) Active Time: $t_{RRD} = t_{RCD} + 1$
- Refresh Cycle Time: $t_{RFC} = t_{RAS} + t_{RP}$
- Exit Self-Refresh Time: $t_{XSR} = t_{RAS} + t_{RP}$
- Load Mode Register to Activate Time: t_{MRD} or t_{RSC} = 3 clock cycles
- Page-Miss Penalty = $t_{RP} + t_{RCD}$

SDRAM Performance

Table 7-1 on page 7-10 lists the data throughput rates for the core or DMA read/write accesses to 32-bit wide SDRAM. For this example, assume all cycles are SCLK cycles and the following SCLK frequency and SDRAM parameters are used:

- SCLK frequency = 133 MHz
- CAS latency = 2 cycles (CL = 2)
- No SDRAM buffering (EBUFE = 0)
- RAS precharge (t_{RP}) = 2 cycles (TRP = 2)
- RAS to CAS delay (t_{RCD}) = 2 cycles (TRCD = 2)
- Active command time (t_{RAS}) = 5 cycles (TRAS = 5)

Bus Request and Grant

The processor can relinquish control of the data and address buses to an external device. The processor three-states its memory interface to allow an external controller to access either external asynchronous or synchronous memory parts.

Operation

When the external device requires access to the bus, it asserts the Bus Request (BR) signal. The BR signal is arbitrated with EAB requests. If no internal request is pending, the external bus request will be granted. The processor initiates a bus grant by:

- Three-stating the data and address buses and the asynchronous memory control signals. The synchronous memory control signals can optionally be three-stated.
- Asserting the Bus Grant (BG) signal.

The processor may halt program execution if the bus is granted to an external device and an instruction fetch or data read/write request is made to external memory. When the external device releases BR, the processor deasserts BG and continues execution from the point at which it stopped.

The processor asserts the $\overline{\text{BGH}}$ pin when it is ready to start another external port access, but is held off because the bus was previously granted.

When the bus has been granted, the BGSTAT bit in the SDSTAT register is set. This bit can be used by the processor to check the bus status to avoid initiating a transaction that would be delayed by the external bus grant.

17 SYSTEM DESIGN

This chapter provides hardware, software and system design information to aid users in developing systems based on the ADSP-BF561 processor. The design options implemented in a system are influenced by cost, performance, and system requirements. In many cases, design issues cited here are discussed in detail in other sections of this manual. In such cases, a reference appears to the corresponding section of the text, instead of repeating the discussion in this chapter.

Pin Descriptions

Refer to the *ADSP-BF561 Embedded Processor Data Sheet* for pin information, including pin numbers for the 256-Lead MBGA package.

Recommendations for Unused Pins

Refer to the *ADSP-BF561 Embedded Processor Data Sheet* for detailed pin descriptions.

Resetting the Processor

In addition to the hardware reset mode provided via the $\overline{\text{RESET}}$ pin, the processor supports several software reset modes. For detailed information on the various modes, see "System Reset and Powerup Configuration" on page 3-12.

The processor state after reset is described in "Reset State" on page 3-10.

Booting the Processor

The processor can be booted via a variety of methods. These include executing from external 16-bit memory, or booting from a ROM configured to load code from 8/16-bit FLASH memory or a serial ROM (8-bit and 16-bit address range). For more information on boot modes, see "Booting Methods" on page 3-19.

Figure 17-1 and Figure 17-2 show the connections necessary for 8-bit and 16-bit parallel booting, respectively. Notice that the address connections are made in the same manner for both 8- and 16-bit peripherals. Only the lower byte of each 16-bit word is accessed if byte-wide memory is used.

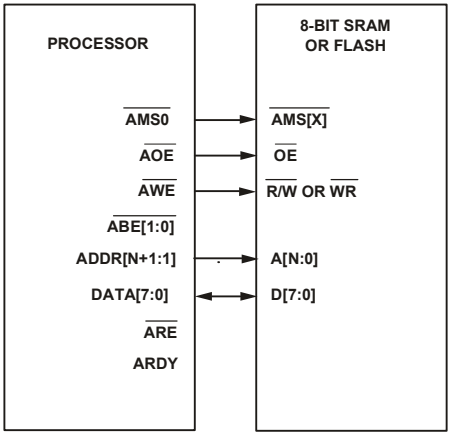

Figure 17-1. Interface to 8-Bit SRAM or Flash

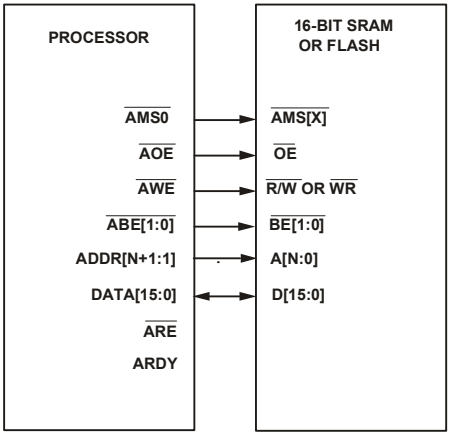

Figure 17-2. Interface to 16-Bit SRAM or Flash

There are no muxes inside the Processor that read the lower 8 bits (connected to external memory) and then move them to the upper 8 bits of the 16-bit word. Therefore, packing into 16-bit words is not supported when connecting 8-bit devices to the external bus. Nevertheless, booting from 8-bit flash is possible because the processor always executes a 16-bit fetch from the external device, regardless of whether an 8-bit or a 16-bit memory device is used. The upper 8 bits are dropped and the lower 8 bits are placed consecutively within internal memory.

Depending upon the choice of 8-bit or 16-bit memories in the Project Options of the VDSP++, the Loader Utility places code and instructions either in the lower 8-bit range of the 16-bit word or over the whole 16-bit range. Therefore it is important to make the choice correctly.

Figure 17-3 and Figure 17-4 show how to connect 8-bit and 16-bit devices to the processor. Note that the address connections are the same.

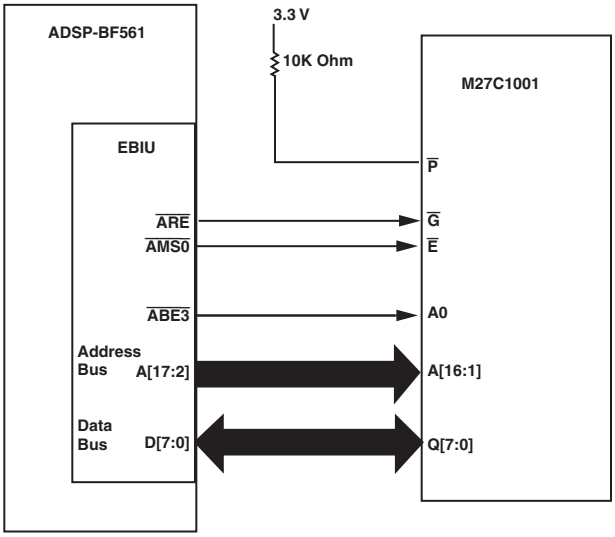

Figure 17-3. ADSP-BF561 to ST EEPROM

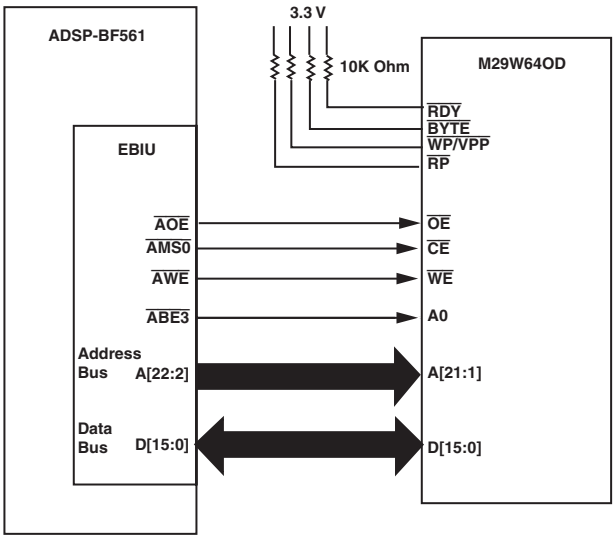

Figure 17-4. ADSP-BF561 to ST FLASH

Managing Clocks

Systems can drive the clock inputs with a crystal oscillator or a buffered, shaped clock derived from an external clock oscillator. The external clock connects to the processor's CLKIN pin. It is not possible to halt, change, or operate CLKIN below the specified frequency during normal operation. The processor uses the clock input (CLKIN) to generate on-chip clocks. These include the core clock (CCLK) and the peripheral clock (SCLK).

Managing Core and System Clocks

The processor produces a (1x-64 for the ADSP-BF561) multiplication of the clock input provided on the CLKIN pin to generate the core clock (CCLK). Additionally, the system clock (SCLK) is derived from the core

clock. This clock is based on a divider ratio that is programmed via the SSEL bit settings. For detailed information about how to set and change CCLK and SCLK frequencies, see "Dynamic Power Management" on page 8-1.

Configuring and Servicing Interrupts

A variety of interrupts are available. They include both core and peripheral interrupts. The processor assigns default core priorities to system-level interrupts. However, these system interrupts can be remapped via the System Interrupt Assignment Registers (SIC_IARx). For more information, see "System Interrupt Assignment Registers (SIC_IARx)" on page 4-30.

The processor core supports nested and non-nested interrupts, as well as self-nested interrupts. For explanations of the various modes of servicing events, please see "Nesting of Interrupts" on page 4-52.

Semaphores

Semaphores provide a mechanism for communication between multiple processors or processes/threads running in the same system. They are used to coordinate resource sharing. For instance, if a process is using a particular resource and another process requires that same resource, it must wait until the first process signals that it is no longer using the resource. This signalling is accomplished via semaphores.

Semaphore coherency is guaranteed by using the Test and Set Byte (Atomic) instruction (TESTSET). The TESTSET instruction performs these functions.

- Loads the half word at memory location pointed to by a P-register. The P-register must be aligned on a half-word boundary.
- Sets CC if the value is equal to zero.
- Stores the value back in its original location (but with the most significant bit (MSB) of the low byte set to 1).

The events triggered by TESTSET are atomic operations. The bus for the memory where the address is located is acquired and not relinquished until the store operation completes. In multithreaded systems, the TESTSET instruction is required to maintain semaphore consistency.

To ensure that the store operation is flushed through any store or write buffers, issue an SSYNC immediately after semaphore release.

The TESTSET instruction can be used to implement binary semaphores or any other type of mutual exclusion method. TESTSET represents a system-level requirement for a multicycle bus lock mechanism.

The processor restricts use of the TESTSET instruction to the external memory region only. Use of the TESTSET instruction to address any other area of the memory map may result in unreliable behavior.

Example Code for Query Semaphore

Listing 17-1 provides an example of a query semaphore that checks the availability of a shared resource.

Listing 17-1. Query Semaphore

```
/* Query semaphore. Denotes “Busy” if its value is non-zero. Wait
until free (or re-schedule thread-- see note below). P0 holds
address of semaphore. */
QUERY:
TESTSET (P0 ) ;
IF !CC JUMP QUERY ;
/* At this point, semaphore has been granted to current thread,
and all other contending threads are postponed because semaphore
value at [P0] is non-zero. Current thread could write thread_id
to semaphore location to indicate current owner of resource. */
R0.L = THREAD_ID ;
B[P0]=R0 ;
/* When done using shared resource, write a zero-byte to [P0] */
R0 = 0 ;
B[P0] = R0 ;
SSYNC ;
/* NOTE: Instead of busy idling in the QUERY loop, one can use an
operating system call to reschedule the current thread. */
```

Data Delays, Latencies and Throughput

For detailed information on latencies and performance estimates on the DMA and External Memory buses, refer to “Chip Bus Hierarchy” on page 7-1.

Bus Priorities

For an explanation of prioritization between the various internal buses, refer to “Chip Bus Hierarchy” on page 7-1.

External Memory Design Issues

This section describes design issues related to external memory including supported SDRAM configurations, examples of SDRAM and SRAM interfaces, and how to avoid bus contention on the external memory bus.

Example Asynchronous Memory Interfaces (16-Bit Memories)

This section shows glueless connections to a 16-bit wide SRAM. Note this example does not require external assertion of ARDY, since the internal wait state counter is sufficient for deterministic access times of memories. Figure 17-5 shows the system interconnect required to support a 16-bit memory.

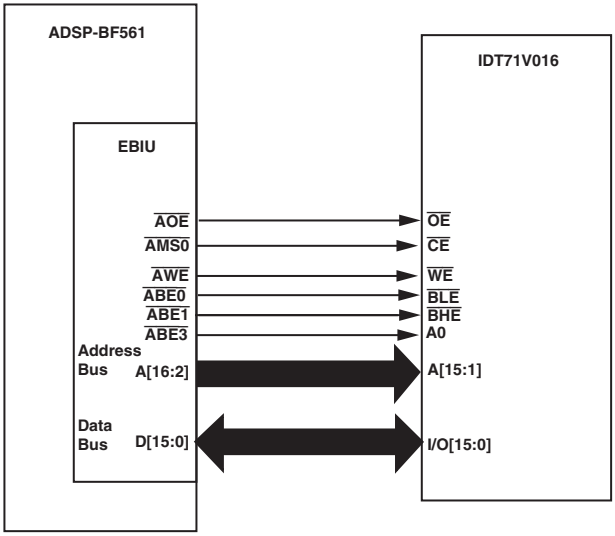

Figure 17-5. Interface to 16-Bit SRAM

Note this application requires that the 16-bit packing mode be *enabled* for this bank of memory. For more information, see "Asynchronous Memory Global Control Register (EBIU_AMGCTL)" on page 16-9. Otherwise, the programming model must ensure that every other 16-bit memory location is accessed starting on an even (byte address[1:0] = 00) 16-bit address.

The byte enable pins can be left open when not supported by the external device. Nevertheless, $\overline{\text{ABE3}}$ needs to be connected to A0 of the external device. For more information, see "Booting the Processor" on page 17-2.

For any 32-bit or 16-bit access to the external device $\overline{\text{ABE2-0}}$ remain low. 8-bit write accesses to the external devices toggle the pins $\overline{\text{ABE1-0}}$ as byte enables.

Example Asynchronous Memory Interfaces (32-Bit Memories)

This section shows glueless connections to 32-bit wide SRAM. Note this interface example does not require external assertion of ARDY, since the internal wait state counter is sufficient for deterministic access times of memories. Figure 17-6 shows the system interconnect required to support 32-bit memories. Note this application requires that the 16-bit packing mode be disabled for this bank of memory. For more information, see "Asynchronous Memory Global Control Register (EBIU_AMGCTL)" on page 16-9. Address pin A2 of the Blackfin Processor must be connected to the pin A0 of the memory.

For any 32-bit access, the $\overline{\text{ABE3-0}}$ pins will remain low. Each 16-bit write access will cause the pins $\overline{\text{ABE1-0}}$ to toggle as a pair and will cause $\overline{\text{ABE3-2}}$ to toggle inverted as a pair. 8-bit write accesses will cause the pins $\overline{\text{ABE3-0}}$ to toggle sequentially.

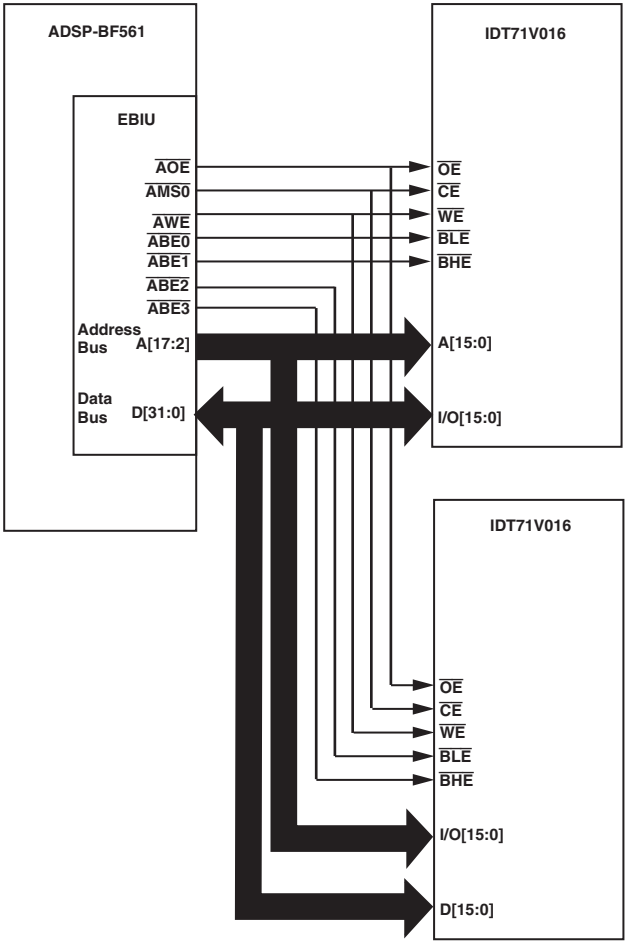

Figure 17-6. ADSP-BF561 to IDT SRAM IDT71V016

High Frequency Design Considerations

Because the processor can operate at very fast clock frequencies, signal integrity and noise problems must be considered for circuit board design

and layout. The following sections discuss these topics and suggest various techniques to use when designing and debugging signal processing systems.

Point-to-Point Connections on Serial Ports

Although the serial ports may be operated at a slow rate, the output drivers still have fast edge rates and for longer distances may require source termination.

You can add a series termination resistor near the pin for point-to-point connections. Typically, serial port applications use this termination method when distances are greater than six inches. For details, see the reference source in "Recommended Reading" on page 17-15 for suggestions on transmission-line termination. Also, see the *ADSP-BF561 Embedded Processor Data Sheet* for output drivers' rise and fall time data.

Signal Integrity

The capacitive loading on high-speed signals should be reduced as much as possible. Loading of buses can be reduced by using a buffer for devices that operate with wait states (for example, DRAMs). This reduces the capacitance on signals tied to the zero-wait-state devices, allowing these signals to switch faster and reducing noise-producing current spikes.

Signal run length (inductance) should also be minimized to reduce ringing. Extra care should be taken with certain signals such as external memory, read, write and acknowledge strobes.

Other recommendations and suggestions to promote signal integrity:

- Use more than one ground plane on the PCB to reduce crosstalk. Be sure to use lots of vias between the ground planes. These planes should be in the center of the PCB.
- Keep critical signals such as clocks, strobes, and bus requests on a signal layer next to a ground plane and away from or laid out perpendicular to other non-critical signals to reduce crosstalk.
- Design for lower transmission line impedances to reduce crosstalk and to allow better control of impedance and delay.
- Experiment with the board and isolate crosstalk and noise issues from reflection issues. This can be done by driving a signal wire from a pulse generator and studying the reflections while other components and signals are passive.

Decoupling Capacitors and Ground Planes

Ground planes must be used for the ground and power supplies. The capacitors should be placed very close to the `VDDEXT` and `VDDINT` pins of the package as shown in Figure 17-7. Use short and fat traces for this. The ground end of the capacitors should be tied directly to the ground plane inside the package footprint of the processor (underneath it, on the bottom of the board), not outside the footprint. A surface-mount capacitor is recommended because of its lower series inductance.

Connect the power plane to the power supply pins directly with minimum trace length. The ground planes must not be densely perforated with vias or traces as their effectiveness is reduced. In addition, there should be several large tantalum capacitors on the board.

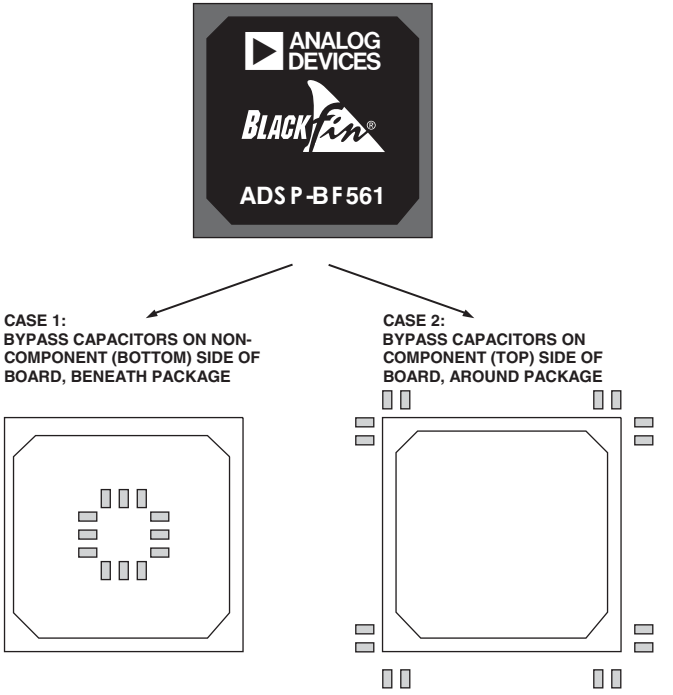

Figure 17-7. Bypass Capacitor Placement

Designs can use either bypass placement, case or combinations of the two. Designs should try to minimize signal feedthroughs that perforate the ground plane.

Oscilloscope Probes

When making high-speed measurements, be sure to use a “bayonet” type or similarly short (< 0.5 inch) ground clip, attached to the tip of the oscilloscope probe. The probe should be a low-capacitance active probe with 3 pF or less of loading. The use of a standard ground clip with 4 inches of

ground lead causes ringing on the displayed trace and makes the signal appear to have excessive overshoot and undershoot. To see the signals accurately, a 1 GHz or better sampling oscilloscope is needed.

Recommended Reading

High-Speed Digital Design: A Handbook of Black Magic, Johnson and Graham, Prentice Hall, Inc., ISBN 0-13-395724-1.

This book is a technical reference that covers the problems encountered in state of the art, high-frequency digital circuit design. It is an excellent source of information and practical ideas. Topics covered in the book include:

- High-speed Properties of Logic Gates
- Measurement Techniques
- Transmission Lines
- Ground Planes and Layer Stacking
- Terminations
- Vias
- Power Systems
- Connectors
- Ribbon Cables
- Clock Distribution
- Clock Oscillators

18 BLACKFIN PROCESSOR DEBUG

The Blackfin processor's debug functionality can be used for software debugging. It also complements some services often found in an operating system (OS) kernel. The functionality is implemented in the processor hardware and is grouped into multiple levels.

A summary of available debug features is shown in Table 18-1.

Table 18-1. Blackfin Debug Features

Debug Feature	Description
Watchpoints	Specify address ranges and conditions that halt the processor when satisfied.
Trace History	Stores the last 16 discontinuous values of the Program Counter in an on-chip trace buffer.
Cycle Count	Provides functionality for all code profiling functions.
Performance Monitoring	Allows internal resources to be monitored and measured non-intrusively.

Watchpoint Unit

By monitoring the addresses on both the instruction bus and the data bus, the Watchpoint Unit provides several mechanisms for examining program behavior. After counting the number of times a particular address is matched, the unit schedules an event based on this count.

In addition, information that the Watchpoint Unit can provide helps optimize code. The unit can also make it easier to maintain executables through code patching.

The Watchpoint Unit contains these MMRs, which are accessible in Supervisor and Emulator modes:

- The Watchpoint Status register (WPSTAT)
- Six Instruction Watchpoint Address registers (WPIA[5:0])
- Six Instruction Watchpoint Address Count registers (WPIACNT[5:0])
- The Instruction Watchpoint Address Control register (WPIACTL)
- Two Data Watchpoint Address registers (WPDA[1:0])
- Two Data Watchpoint Address Count registers (WPDACNT[1:0])
- The Data Watchpoint Address Control register (WPDACTL)

Two operations implement instruction watchpoints:

- The values in the six Instruction Watchpoint Address registers, WPIA[5:0], are compared to the address on the instruction bus.
- Corresponding count values in the Instruction Watchpoint Address Count registers, WPIACNT[5:0], are decremented on each match.

The six Instruction Watchpoint Address registers may be further grouped into three ranges of instruction-address-range watchpoints. The ranges are identified by the addresses in WPIA0 to WPIA1, WPIA2 to WPIA3, and WPIA4 to WPIA5.

The address ranges stored in WPIA0, WPIA1, WPIA2, WPIA3, WPIA4 and WPIA5 must satisfy these conditions:

- WPIA0 <= WPIA1
- WPIA2 <= WPIA3
- WPIA4 <= WPIA5

Two operations implement data watchpoints:

- The values in the two Data Watchpoint Address registers, WPDA[1:0], are compared to the address on the data buses.
- Corresponding count values in the Data Watchpoint Address Count registers, WPDACNT[1:0], are decremented on each match.

The two Data Watchpoint Address registers may be further grouped together into one data-address-range watchpoint, WPDA[1:0].

The instruction and data count value registers must be loaded with one less than the number of times the watchpoint must match. After the count value reaches zero, the subsequent watchpoint match results in an exception or emulation event.

Note count values must be reinitialized after the event has occurred.

An event can also be triggered on a combination of the instruction and data watchpoints. If the WPAND bit in the WPIACTL register is set, then an event is triggered only when both an instruction address watchpoint matches *and* a data address watchpoint matches. If the WPAND bit is 0, then an event is triggered when any of the enabled watchpoints or watchpoint ranges match.

To enable the Watchpoint Unit, the WPPWR bit in the WPIACTL register must be set. If WPPWR=1, then the individual watchpoints and watchpoint

ranges may be enabled using the specific enable bits in the WPIACTL and WPDACTL MMRs. If WPPWR=0, then all watchpoint activity is disabled.

Instruction Watchpoints

Each instruction watchpoint is controlled by three bits in the WPIACTL register, as shown in Table 18-2.

Table 18-2. WPIACTL Control Bits

Bit Names	Description
EMUSWx	Determines whether an instruction-address match causes either an emulation event or an exception event.
WPICNTENx	Enables the 16-bit counter that counts the number of address matches. If the counter is disabled, then every match causes an event.
WPIAENx	Enables the address watchpoint activity.

When two watchpoints are associated to form a range, two additional bits are used, as shown in Table 18-3.

Table 18-3. WPIACTL Watchpoint Range Control Bits

Bit Names	Description
WPIRENxy	Indicates the two watchpoints that are to be associated to form a range.
WPIRINVxy	Determines whether an event is caused by an address within the range identified or outside of the range identified.

Code patching allows software to replace sections of existing code with new code. The watchpoint registers are used to trigger an exception at the

start addresses of the earlier code. The exception routine then vectors to the location in memory that contains the new code.

On the processor, code patching can be achieved by writing the start address of the earlier code to one of the WPIAn registers and setting the corresponding EMUSWx bit to trigger an exception. In the exception service routine, the WPSTAT register is read to determine which watchpoint triggered the exception. Next, the code writes the start address of the new code in the RETX register and then returns from the exception to the new code. Because the exception mechanism is used for code patching, event service routines of the same or higher priority (exception, NMI, and reset routines) cannot be patched.

A write to the WPSTAT MMR clears all the sticky status bits. The data value written is ignored.

Instruction Watchpoint Address Registers (WPIAn)

When the Watchpoint Unit is enabled, the values in the Instruction Watchpoint Address (WPIAn) registers are compared to the address on the instruction bus. Corresponding count values in the Instruction Watchpoint Address Count registers (WPIACNTn) are decremented on each match.

Figure 18-1 shows the Instruction Watchpoint Address registers, WPIA[5:0].

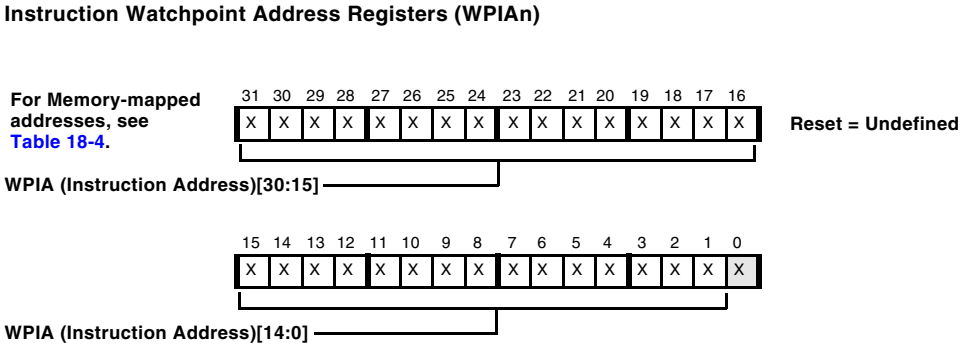

Figure 18-1. Instruction Watchpoint Address Registers

Table 18-4. Instruction Watchpoint Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
WPIA0	0xFFE0 7040
WPIA1	0xFFE0 7044
WPIA2	0xFFE0 7048
WPIA3	0xFFE0 704C
WPIA4	0xFFE0 7050
WPIA5	0xFFE0 7054

Instruction Watchpoint Address Count Registers (WPIACNTn)

When the Watchpoint Unit is enabled, the count values in the Instruction Watchpoint Address Count registers (`WPIACNT[5:0]`) are decremented each time the address or the address bus matches a value in the `WPIAn` registers. Load the `WPIACNTn` register with a value that is one less than the number of times the watchpoint must match before triggering an event

(see Figure 18-2). The WPIACNTn register will decrement to 0x0000 when the programmed count expires.

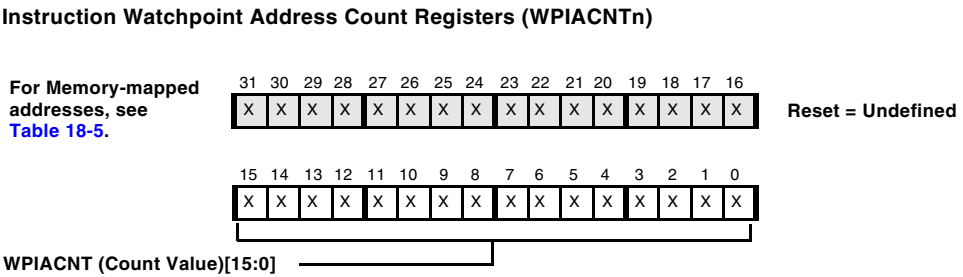

Figure 18-2. Instruction Watchpoint Address Count Register

Table 18-5. Instruction Watchpoint Address Count Register Memory-Mapped Addresses

Register Name	Memory-Mapped Address
WPIACNT0	0xFFE0 7080
WPIACNT1	0xFFE0 7084
WPIACNT2	0xFFE0 7088
WPIACNT3	0xFFE0 708C
WPIACNT4	0xFFE0 7090
WPIACNT5	0xFFE0 7094

Instruction Watchpoint Address Control Register (WPIACTL)

Three bits in the Instruction Watchpoint Address Control register (WPIACTL) control each instruction watchpoint. Figure 18-3 describes the upper half of the register. Figure 18-4 describes the lower half of the

register. For more information about the bits in this register, see "Instruction Watchpoints" on page 18-4.

The bits in the WPIACTL register have no effect unless the WPPWR bit is set.

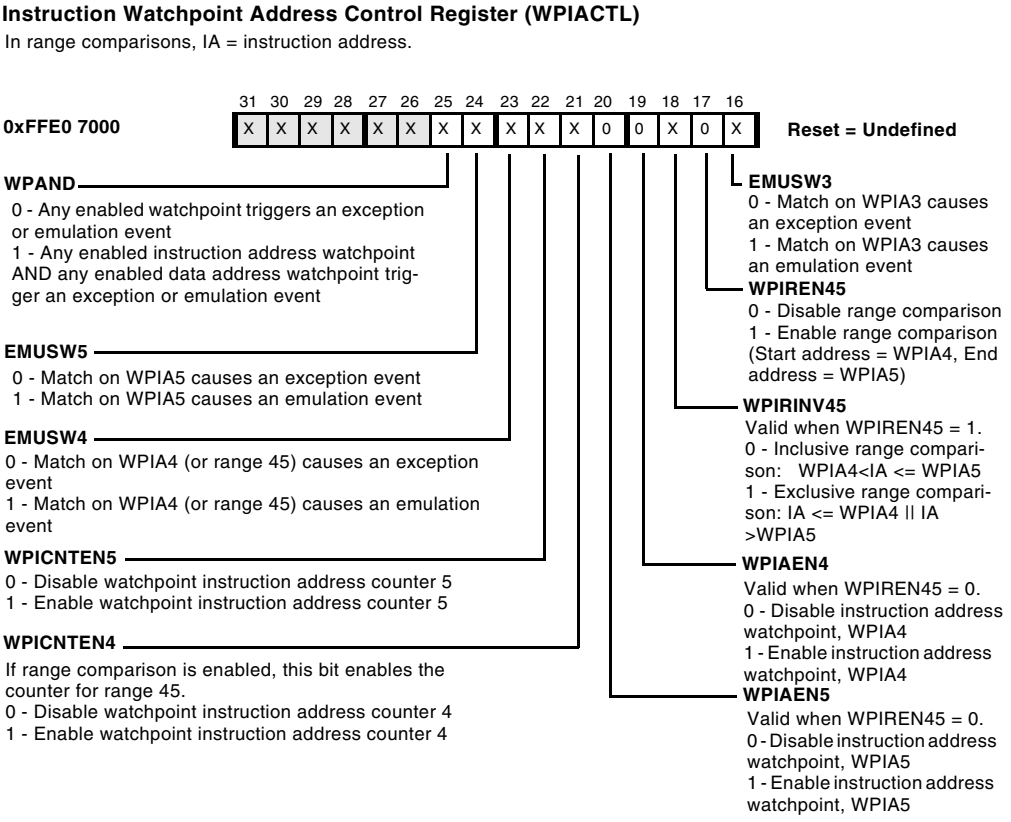

Figure 18-3. Instruction Watchpoint Address Control Register (WPIACTL)[31:16]

Instruction Watchpoint Address Control Register (WPIACTL)

In range comparisons, IA = instruction address.

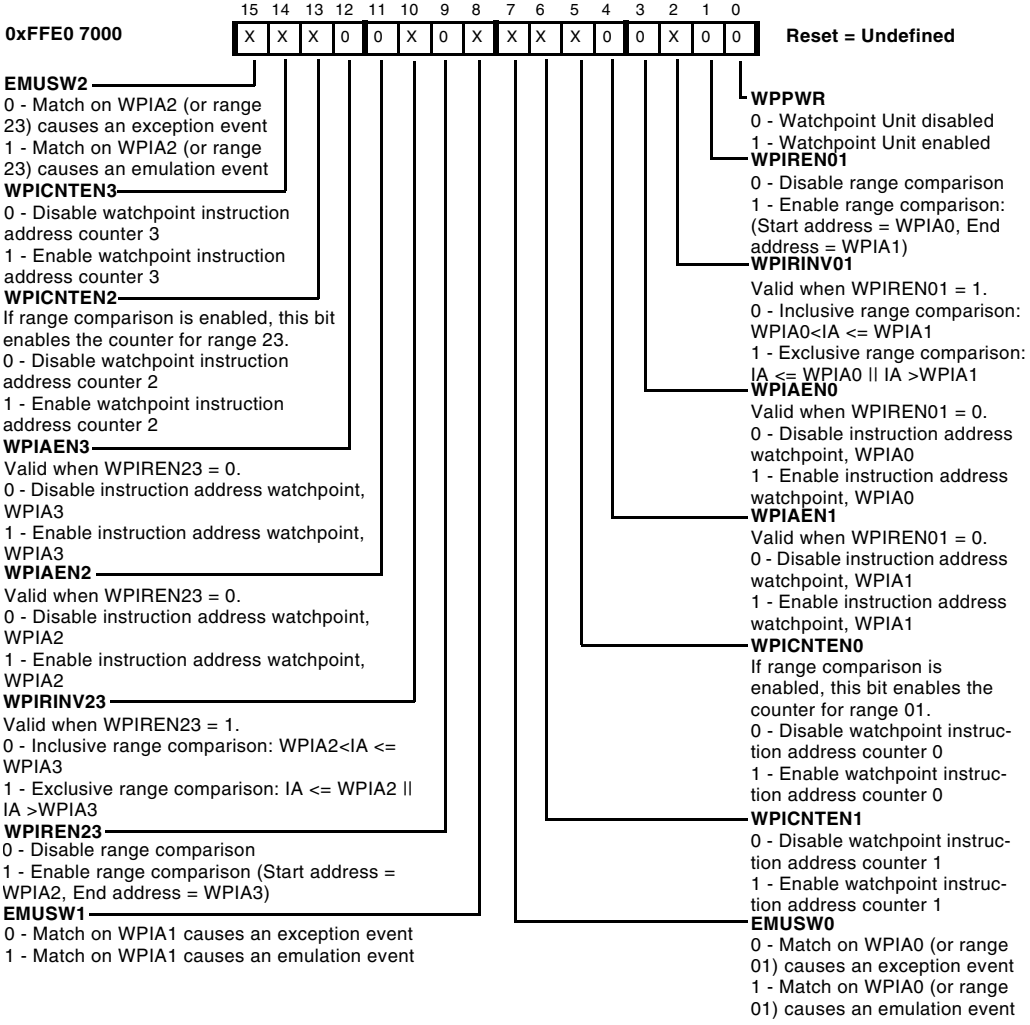

Figure 18-4. Instruction Watchpoint Address Control Register (WPIACTL)[15:0]

Data Address Watchpoints

Each data watchpoint is controlled by four bits in the WPDACTL register, as shown in Table 18-6.

Table 18-6. Data Address Watchpoints

Bit Names	Description
WPDACCn	Determine whether the match should be on a read or write access.
WPDSRCn	Determine which DAG the unit should monitor.
WPDCNTENn	Enable the counter that counts the number of address matches. If the counter is disabled, then every match causes an event.
WPDAENn	Enable the data watchpoint activity.

When the two watchpoints are associated to form a range, two additional bits are used, as shown in Table 18-7.

Table 18-7. WPDACTL Watchpoint Control Bits

Bit Name	Description
WPDREN01	Indicates the two watchpoints that are to be associated to form a range.
WPDRINV01	Determines whether an event is caused by an address within the range identified or outside the range.

Note Data address watchpoints always trigger emulation events.

Data Watchpoint Address Registers (WPDAn)

When the Watchpoint Unit is enabled, the values in the Data Watchpoint Address (WPDAn) registers are compared to the address on the data buses. Corresponding count values in the Data Watchpoint Address Count registers (WPDACNTn) are decremented on each match.

Figure 18-5 shows the Data Watchpoint Address registers, WPDA[1:0].

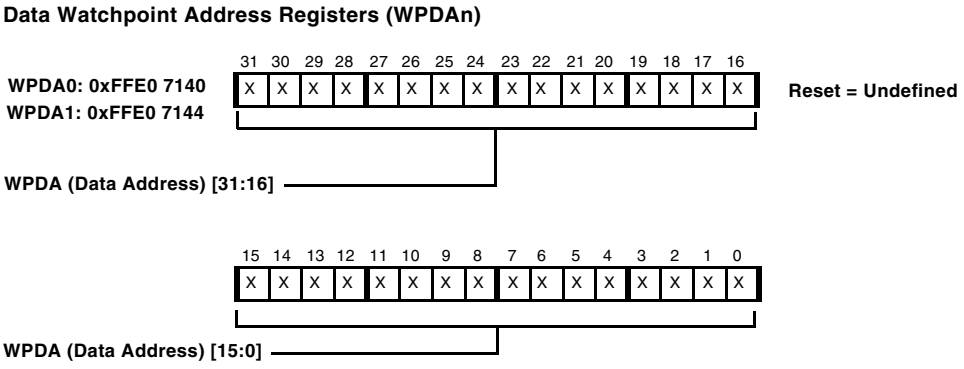

Figure 18-5. Data Watchpoint Address Registers

Data Watchpoint Address Count Value Registers (WPDACNTn)

When the Watchpoint Unit is enabled, the count values in the Data Watchpoint Address Count Value (WPDACNTn) registers are decremented each time the address or the address bus matches a value in the WPDAn registers. Load this WPDACNTn register with a value that is one less than the number of times the watchpoint must match before triggering an event. The WPDACNTn register will decrement to 0x0000 when the programmed

count expires. Figure 18-6 shows the Data Watchpoint Address Count Value registers, WPDACNT[1:0].

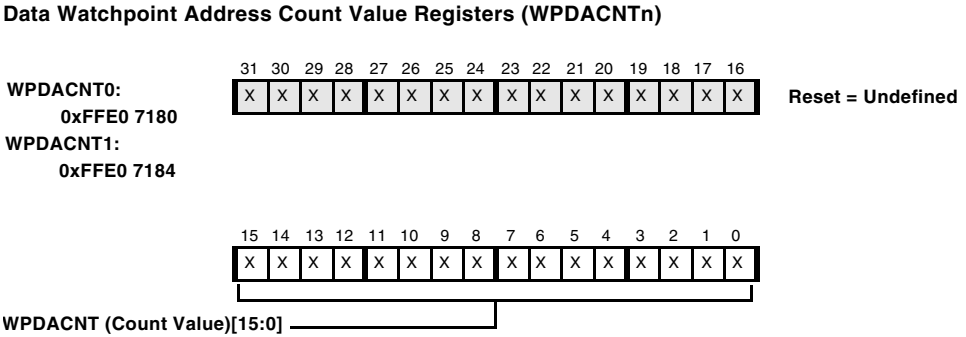

Figure 18-6. Data Watchpoint Address Count Value Registers

Data Watchpoint Address Control Register (WPDACTL)

For more information about the bits in the Data Watchpoint Address Control (WPDACTL) register, see "Data Address Watchpoints" on page 18-10 and Figure 18-7.

Watchpoint Status Register (WPSTAT)

The Watchpoint Status (WPSTAT) register (Figure 18-8) monitors the status of the watchpoints. It may be read and written in Supervisor or Emulator modes only. When a watchpoint or watchpoint range matches, this register reflects the source of the watchpoint. The status bits in the WPSTAT register are sticky, and all of them are cleared when any write, regardless of the value, is performed to the register.

Data Watchpoint Address Control Register (WPDACTL)

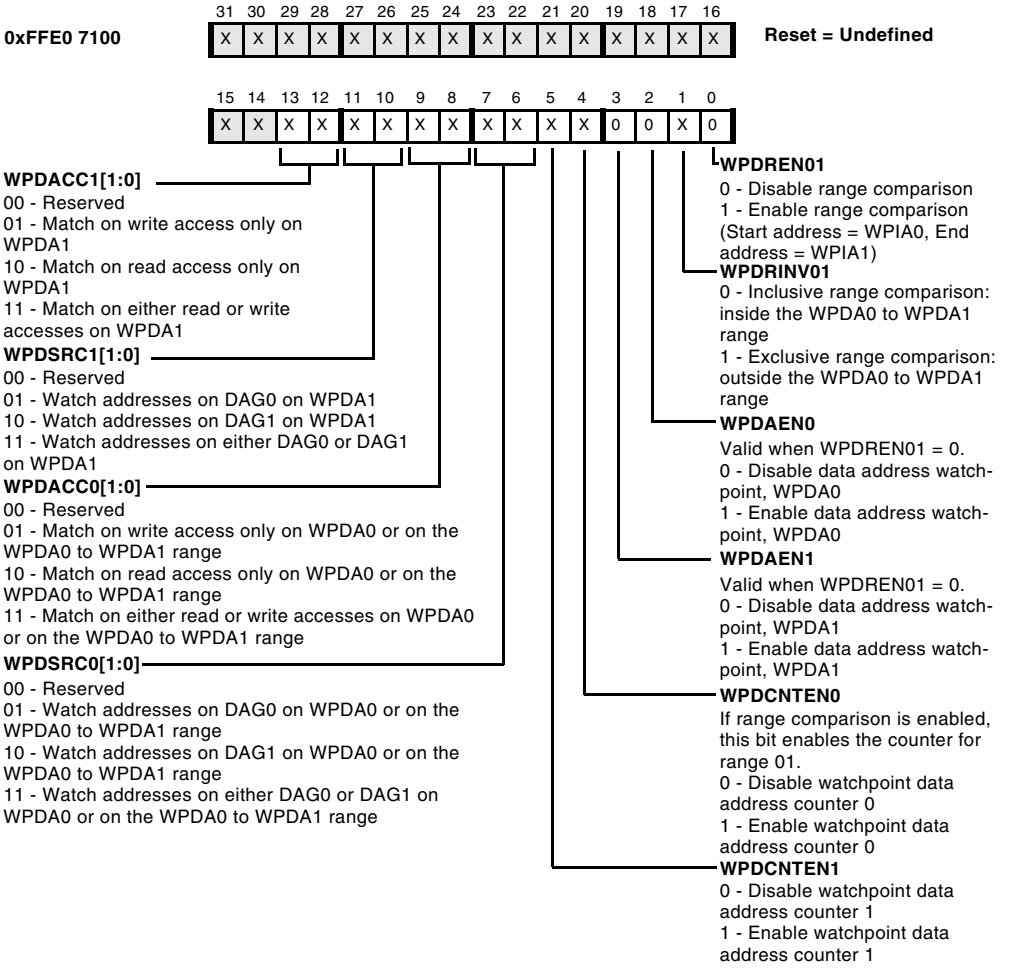

Figure 18-7. Data Watchpoint Address Control Register

Watchpoint Status Register (WPSTAT)

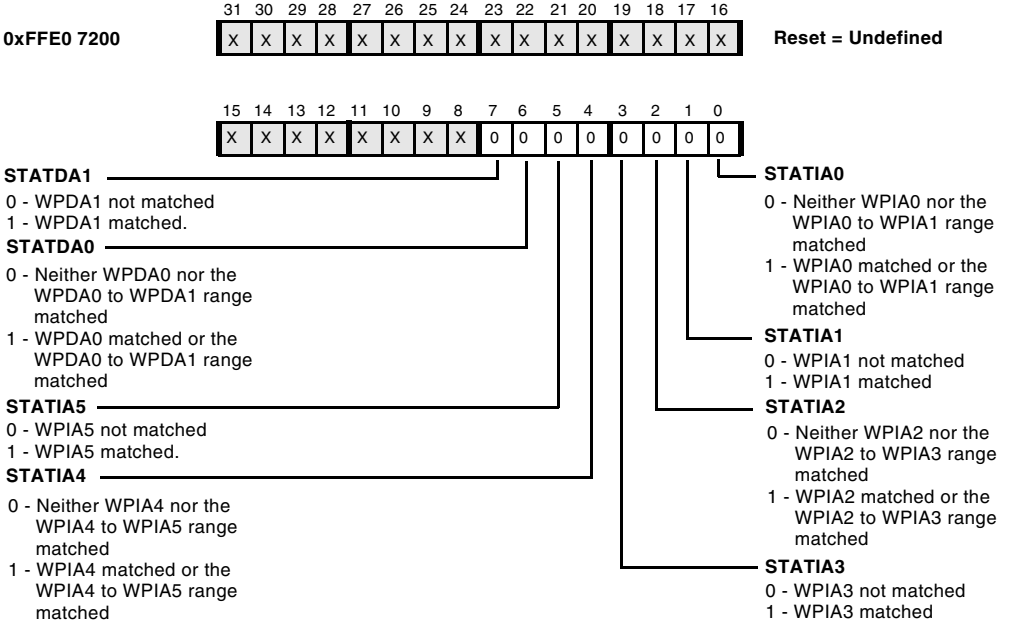

Figure 18-8. Watchpoint Status Register

Trace Unit

The Trace Unit stores a history of the last 16 changes in program flow taken by the program sequencer. The history allows the user to recreate the program sequencer's recent path.

The trace buffer can be enabled to cause an exception when full. The exception service routine associated with the exception saves trace buffer entries to memory. Thus, the complete path of the program sequencer since the trace buffer was enabled can be recreated.

Changes in program flow because of zero-overhead loops are not stored in the trace buffer. For debugging code that is halted within a zero-overhead loop, the iteration count is available in the loop count registers, LC0 and LC1.

The trace buffer can be configured to omit recording of changes in program flow that match either the last entry or one of the last two entries. Omitting one of these entries from the record prevents the trace buffer from overflowing because of loops in the program. Because zero-overhead loops are not recorded in the trace buffer, this feature can be used to prevent trace overflow from loops that are nested four deep.

When read, the Trace Buffer register (TBUF) returns the top value from the Trace Unit stack, which contains as many as sixteen entries. Each entry contains a pair of branch source and branch target addresses. A read of TBUF returns the newest entry first, starting with the branch destination. The next read provides the branch source address.

The number of valid entries in TBUF is held in the TBUFCNT field of the TBUFSTAT register. On every second read, TBUFCNT is decremented. Because each entry corresponds to two pieces of data, a total of 2 × TBUFCNT reads empties TBUF.

Discontinuities that are the same as either of the last two entries in the trace buffer are not recorded.

Because reading the trace buffer is a destructive operation, it is recommended that TBUF be read in a non-interruptible section of code.

Note if single-level compression has occurred, the LSB of the branch target address is set. If two-level compression has occurred, the LSB of the branch source address is set.

Trace Buffer Control Register (TBUFCTL)

The Trace Unit is enabled by two control bits in the Trace Buffer Control register (TBUFCTL) register. First, the Trace Unit must be activated by setting the TBUFPWR bit. If TBUFPWR = 1, then setting TBUFEN to 1 enables the Trace Unit.

Figure 18-9 describes the Trace Buffer Control register (TBUFCTL). If TBUFOVF = 1, then the Trace Unit does not record discontinuities in the exception, NMI, and reset routines.

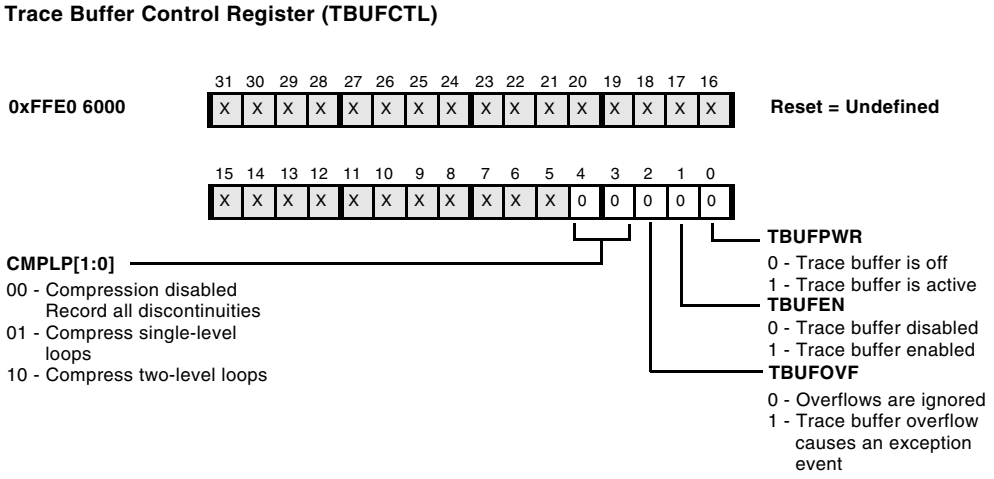

Figure 18-9. Trace Buffer Control Register

Trace Buffer Status Register (TBUFSTAT)

Figure 18-10 shows the Trace Buffer Status register (TBUFSTAT). Two reads from TBUF decrements TBUFCNT by 1.

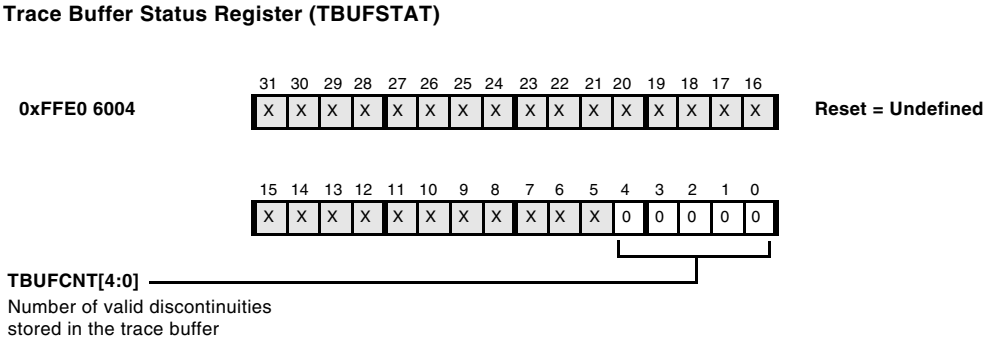

Figure 18-10. Trace Buffer Status Register

Trace Buffer Register (TBUF)

Figure 18-11 shows the Trace Buffer register (TBUF). The first read returns the latest branch target address. The second read returns the latest branch source address.

The Trace Unit does not record changes in program flow in:

- Emulator mode
- The exception or higher priority service routines (if TBUFOVF = 1)
- In the exception service routine, the program flow discontinuities may be read from TBUF and stored in memory by using the code shown in Listing 18-1.

While TBUF is being read, be sure to disable the trace buffer from recording new discontinuities.

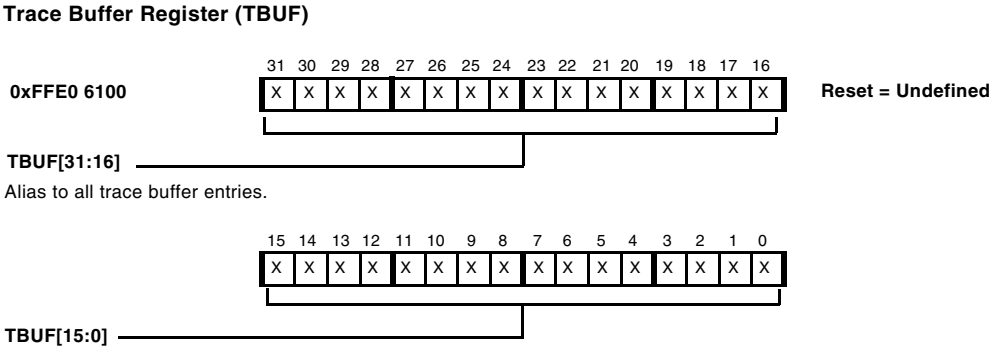

Figure 18-11. Trace Buffer Register

Code to Recreate the Execution Trace in Memory

Listing 18-1 provides code that recreates the entire execution trace in memory.

Listing 18-1. Recreating the Execution Trace in Memory

```
[--sp] = (r7:7, p5:2); /* save registers used in this routine */
p5 = 32; /* 32 reads are needed to empty TBUF */
p2.l = buf; /* pointer to the header (first location) of the
software trace buffer */
p2.h = buf; /* the header stores the first available empty buf
location for subsequent trace dumps */

p4 = [p2++]; /* get the first available empty buf location from
the buf header */
p3.l = TBUF & 0xffff; /* low 16-bits of TBUF */
p3.h = TBUF >> 16; /* high 16-bits of TBUF */

lsetup(loop1_start, loop1_end) lc0 = p5;
loop1_start: r7 = [p3]; /* read from TBUF */
loop1_end: [p4++] = r7; /* write to memory and increment */
```

```
[p2] = p4; /* pointer to the next available buf location is saved
in the header of buf */
(r7:7, p5:3) = [sp++]; /* restore saved registers */
```

Performance Monitoring Unit

Two 32-bit counters, the Performance Monitor Counter registers (PFCNTR[1:0]) and the Performance Control register (PFCTL), count the number of occurrences of an event from within a processor core unit during a performance monitoring period. These registers provide feedback indicating the measure of load balancing between the various resources on the chip so that expected and actual usage can be compared and analyzed. In addition, events such as mispredictions and hold cycles can also be monitored.

Performance Monitor Counter Registers (PFCNTRn)

Figure 18-12 shows the Performance Monitor Counter registers, PFCNTR[1:0]. The PFCNTR0 register contains the count value of performance counter 0. The PFCNTR1 register contains the count value of performance counter 1.

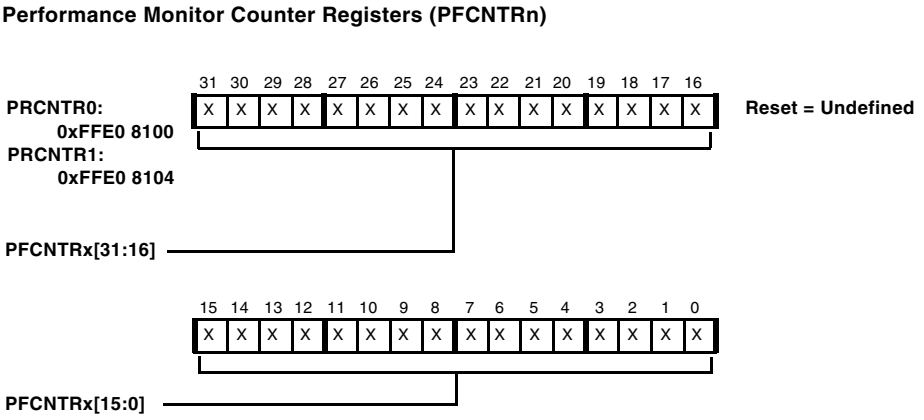

Figure 18-12. Performance Monitor Counter Registers

Performance Monitor Control Register (PFCTL)

To enable the Performance Monitoring Unit, set the PFPWR bit in the Performance Monitor Control (PFCTL) register (see Figure 18-13). Once the unit is enabled, individual count-enable bits (PFCENn) take effect. Use the PFCENx bits to enable or disable the performance monitors in User mode, Supervisor mode, or both. Use the PEMUSWx bits to select the type of event triggered.

Performance Monitor Control Register (PFCTL)

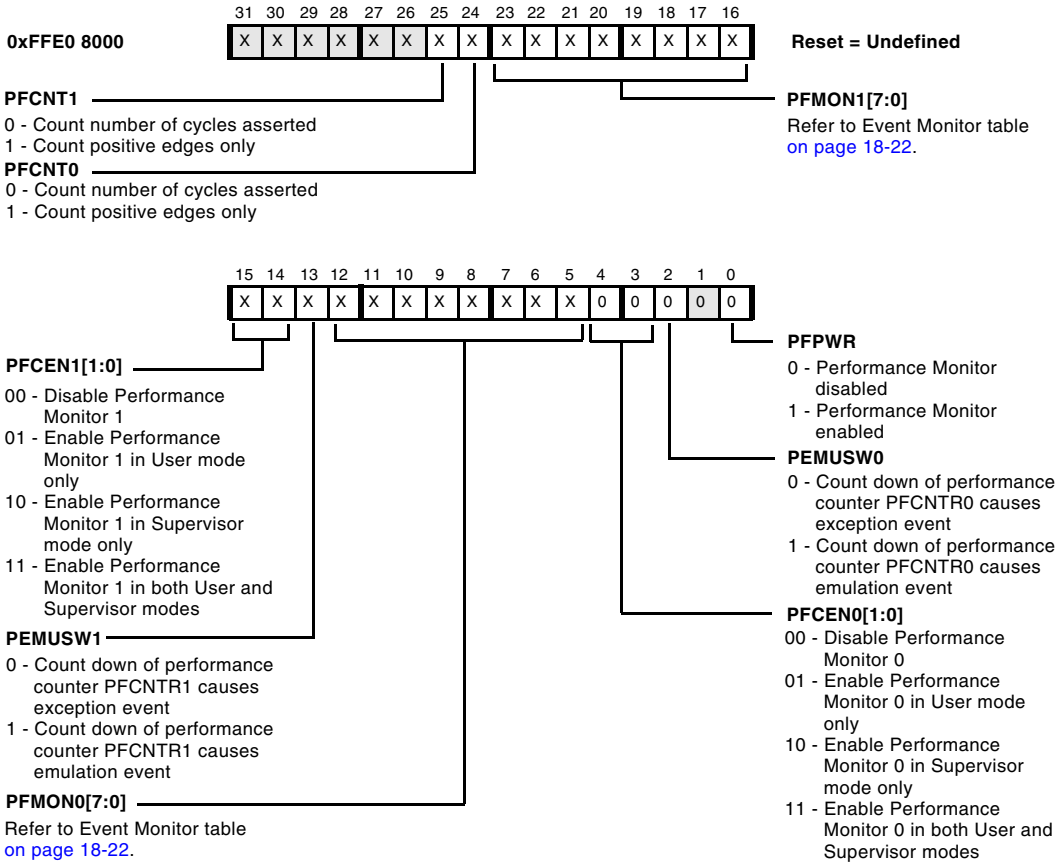

Figure 18-13. Performance Monitor Control Register

Event Monitor Table

Table 18-8 identifies events that cause the Performance Monitor Counter registers (PFMON0 or PFMON1) to increment.

Table 18-8. Event Monitor Table

PFMONx Fields	Events That Cause the Count Value to Increment
0x00	Loop 0 iterations
0x01	Loop 1 iterations
0x02	Loop buffer 0 not optimized
0x03	Loop buffer 1 not optimized
0x04	PC invariant branches. Requires Trace Buffer to be enabled - see TBUFCTL on page 18-16.
0x06	Conditional branches
0x09	Total branches (calls, returns, branches, but not interrupts). Requires Trace Buffer to be enabled - see TBUFCTL on page 18-16.
0x0A	Stalls due to CSYNC, SSYNC
0x0B	EXCPT instructions
0x0C	CSYNC, SSYNC instructions
0x0D	Committed instructions
0x0E	Interrupts taken
0x0F	Misaligned address violation exceptions
0x10	Stall cycles due to read after write hazards on DAG registers
0x13	Stall cycles due to RAW data hazards in computes
0x80	Code memory fetches are postponed due to DMA collisions (minimum count of two per event)
0x81	Code memory TAG stalls (cache misses, or FlushI operations, count of 3 per FlushI). Note code memory stall results in a processor stall only if instruction assembly unit FIFO empties.

Table 18-8. Event Monitor Table (Cont'd)

PFMONx Fields	Events That Cause the Count Value to Increment
0x82	Code memory fill stalls (cacheable or non-cacheable). Note code memory stall results in a processor stall only if instruction assembly unit FIFO empties.
0x83	Code memory 64 bit words delivered to processor instruction assembly unit
0x90	Processor stalls to memory
0x91	Data memory stalls to processor are not hidden by processor stall
0x92	Data memory store buffer full stalls
0x93	Data memory write buffer full stalls due to high to low priority code transition
0x94	Data memory store buffer forward stalls due to lack of committed data from processor
0x95	Data memory fill buffer stalls
0x96	Data memory array or TAG collision stalls (DAG to DAG, or DMA to DAG)
0x97	Data memory array collision stalls (DAG to DAG or DMA to DAG)
0x98	Data memory stalls
0x99	Data memory stalls sent to processor
0x9A	Data memory cache fills completed to Bank A
0x9B	Data memory cache fills completed to Bank B
0x9C	Data memory cache victims delivered from Bank A
0x9D	Data memory cache victims delivered from Bank B
0x9E	Data memory cache high priority fills requested
0x9F	Data memory cache low priority fills requested

Cycle Counter

The cycle counter counts CCLK cycles while the program is executing. All cycles, including execution, wait state, interrupts, and events, are counted while the processor is in User or Supervisor mode, but the cycle counter stops counting in Emulator mode.

The cycle counter is 64 bits wide and increments every cycle. The count value is stored in two 32-bit registers, CYCLES and CYCLES2. The least significant 32 bits (LSBs) are stored in CYCLES. The most significant 32 bits (MSBs) are stored in CYCLES2.

To ensure read coherency, first read CYCLES, then CYCLES2, and then CYCLES again, to detect if an overflow has occurred in the LSBs during the read operations.

To enable the cycle counters, set the CCEN bit in the SYSCFG register. The following example shows how to use the cycle counter:

```
R2 = 0;
CYCLES = R2;
CYCLES2 = R2;
R2 = SYSCFG;
BITSET(R2,1);
SYSCFG = R2;
/* Insert code to be benchmarked here. */
R2 = SYSCFG;
BITCLR(R2,1);
SYSCFG = R2;
```

Execution Cycle Count Registers (CYCLES and CYCLES2)

The Execution Cycle Count registers CYCLES and CYCLES2 are shown in Figure 18-14. This 64-bit counter increments every CCLK cycle. CYCLES contains the least significant 32 bits of the cycle counter's 64-bit count value. The most significant 32 bits are contained by CYCLES2.

The CYCLES and CYCLES2 registers are not system MMRs, but are instead system registers. A full listing of system registers (Sysreg) can be found on page 2-8.

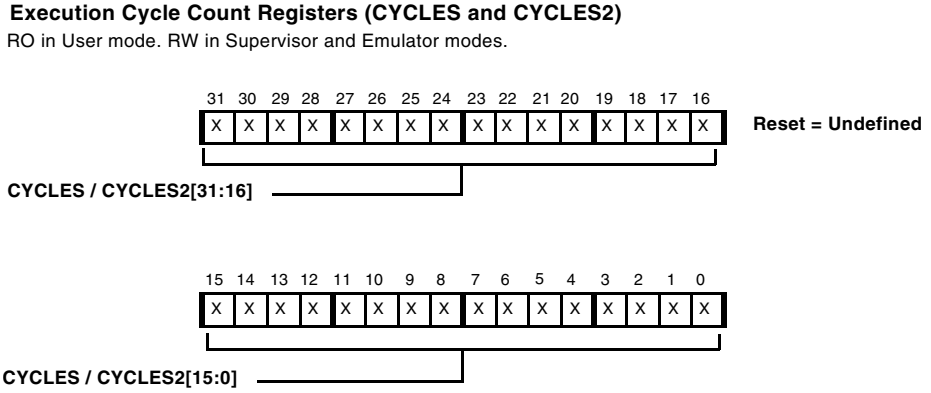

Figure 18-14. Execution Cycle Count Registers

Product Identification Registers

The 32-bit DSP Device ID register (`DSPID`) is a core MMR that contains core identification and revision fields for the core.

DSP Device ID Register (DSPID)

The DSP Device ID register (DSPID), shown in Figure 18-15, is read-only and part of the core.

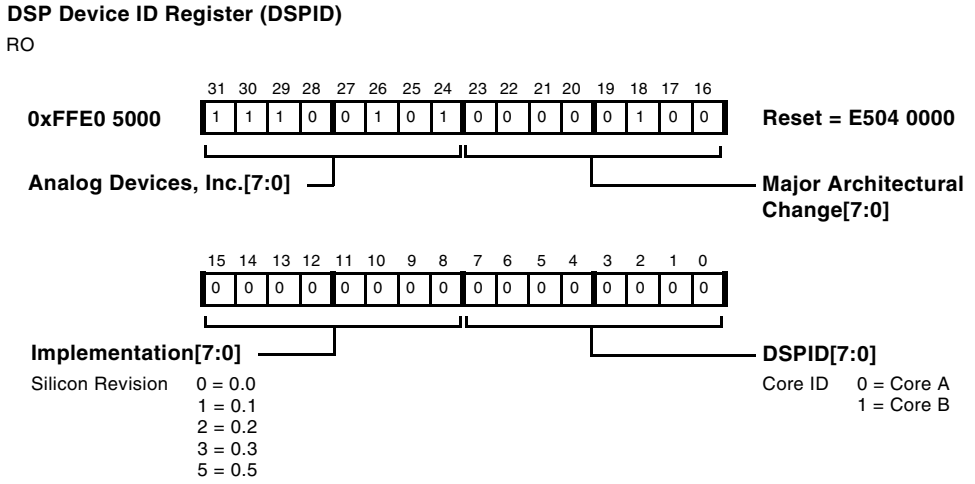

Figure 18-15. DSP Device ID Register

A BLACKFIN PROCESSOR CORE MMR ASSIGNMENTS

The Blackfin processor's memory-mapped registers (MMRs) are in the address range 0xFFE0 0000 – 0xFFFF FFFF.

All core MMRs are 32 bits and can be accessed only with 32-bit aligned memory operations.

Each core has its own set of core MMRs with identical mapping. The following lists apply to both cores.

This appendix lists core MMR addresses and register names. To find more information about an MMR, refer to the page shown in the "See Section" column. When viewing the PDF version of this document, click a reference in the "See Section" column to jump to additional information about the MMR.

L1 Data Memory Controller Registers

L1 Data Memory Controller registers (0xFFE0 0000 – 0xFFE0 0404) are shown in Table A-1.

Table A-1. L1 Data Memory Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 0004	DMEM_CONTROL	"Data Memory Control Register (DMEM_CONTROL)" on page 6-24
0xFFE0 0008	DCPLB_STATUS	"DCPLB and ICPLB Status Registers (DCPLB_STATUS, ICPLB_STATUS)" on page 6-62
0xFFE0 000C	DCPLB_FAULT_ADDR	"DCPLB and ICPLB Fault Address Registers (DCPLB_FAULT_ADDR, ICPLB_FAULT_ADDR)" on page 6-64
0xFFE0 0100	DCPLB_ADDR0	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0104	DCPLB_ADDR1	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0108	DCPLB_ADDR2	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 010C	DCPLB_ADDR3	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0110	DCPLB_ADDR4	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0114	DCPLB_ADDR5	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0118	DCPLB_ADDR6	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 011C	DCPLB_ADDR7	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0120	DCPLB_ADDR8	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60

Table A-1. L1 Data Memory Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 0124	DCPLB_ADDR9	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0128	DCPLB_ADDR10	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 012C	DCPLB_ADDR11	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0130	DCPLB_ADDR12	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0134	DCPLB_ADDR13	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0138	DCPLB_ADDR14	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 013C	DCPLB_ADDR15	"DCPLB Address Registers (DCPLB_ADDRx)" on page 6-60
0xFFE0 0200	DCPLB_DATA0	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0204	DCPLB_DATA1	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0 xFFE0 0208	DCPLB_DATA2	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 020C	DCPLB_DATA3	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0210	DCPLB_DATA4	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0214	DCPLB_DATA5	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0218	DCPLB_DATA6	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 021C	DCPLB_DATA7	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57

Table A-1. L1 Data Memory Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 0220	DCPLB_DATA8	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0224	DCPLB_DATA9	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0228	DCPLB_DATA10	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 022C	DCPLB_DATA11	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0230	DCPLB_DATA12	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0234	DCPLB_DATA13	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0238	DCPLB_DATA14	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 023C	DCPLB_DATA15	"DCPLB Data Registers (DCPLB_DATAx)" on page 6-57
0xFFE0 0300	DTEST_COMMAND	"Data Test Command Register (DTEST_COMMAND)" on page 6-41
0xFFE0 0400	DTEST_DATA0	"Data Test Data 0 Register (DTEST_DATA0)" on page 6-42
0xFFE0 0404	DTEST_DATA1	"Data Test Data 1 Register (DTEST_DATA1)" on page 6-42

L1 Instruction Memory Controller Registers

L1 Instruction Memory Controller registers (0xFFE0 1004 – 0xFFE0 1404) are shown in Table A-2.

Table A-2. L1 Instruction Memory Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 1004	IMEM_CONTROL	"Instruction Memory Control Register (IMEM_CONTROL)" on page 6-7
0xFFE0 1008	ICPLB_STATUS	"DCPLB and ICPLB Status Registers (DCPLB_STATUS, ICPLB_STATUS)" on page 6-62
0xFFE0 100C	ICPLB_FAULT_ADDR	"DCPLB and ICPLB Fault Address Registers (DCPLB_FAULT_ADDR, ICPLB_FAULT_ADDR)" on page 6-64
0xFFE0 1100	ICPLB_ADDR0	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1104	ICPLB_ADDR1	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1108	ICPLB_ADDR2	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 110C	ICPLB_ADDR3	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1110	ICPLB_ADDR4	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1114	ICPLB_ADDR5	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1118	ICPLB_ADDR6	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 111C	ICPLB_ADDR7	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61

Table A-2. L1 Instruction Memory Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 1120	ICPLB_ADDR8	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1124	ICPLB_ADDR9	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1128	ICPLB_ADDR10	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 112C	ICPLB_ADDR11	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1130	ICPLB_ADDR12	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1134	ICPLB_ADDR13	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1138	ICPLB_ADDR14	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 113C	ICPLB_ADDR15	"ICPLB Address Registers (ICPLB_ADDRx)" on page 6-61
0xFFE0 1200	ICPLB_DATA0	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1204	ICPLB_DATA1	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1208	ICPLB_DATA2	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 120C	ICPLB_DATA3	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1210	ICPLB_DATA4	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1214	ICPLB_DATA5	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1218	ICPLB_DATA6	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56

Table A-2. L1 Instruction Memory Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 121C	ICPLB_DATA7	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1220	ICPLB_DATA8	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1224	ICPLB_DATA9	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1228	ICPLB_DATA10	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 122C	ICPLB_DATA11	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1230	ICPLB_DATA12	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1234	ICPLB_DATA13	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1238	ICPLB_DATA14	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 123C	ICPLB_DATA15	"ICPLB Data Registers (ICPLB_DATAx)" on page 6-56
0xFFE0 1300	ITEST_COMMAND	"ITEST_COMMAND Register" on page 6-21
0XFFE0 1400	ITEST_DATA0	"ITEST_DATA0 Register" on page 6-23
0XFFE0 1404	ITEST_DATA1	"ITEST_DATA1 Register" on page 6-22

Interrupt Controller Registers

Interrupt Controller registers (0xFFE0 2000 – 0xFFE0 2110) are shown in Table A-3.

Table A-3. Interrupt Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 2000	EVT0 (EMU)	"Core Event Vector Table" on page 4-40
0xFFE0 2004	EVT1 (RST)	"Core Event Vector Table" on page 4-40
0xFFE0 2008	EVT2 (NMI)	"Core Event Vector Table" on page 4-40
0xFFE0 200C	EVT3 (EVX)	"Core Event Vector Table" on page 4-40
0xFFE0 2010	EVT4	"Core Event Vector Table" on page 4-40
0xFFE0 2014	EVT5 (IVHW)	"Core Event Vector Table" on page 4-40
0xFFE0 2018	EVT6 (TMR)	"Core Event Vector Table" on page 4-40
0xFFE0 201C	EVT7 (IVG7)	"Core Event Vector Table" on page 4-40
0xFFE0 2020	EVT8 (IVG8)	"Core Event Vector Table" on page 4-40
0xFFE0 2024	EVT9 (IVG9)	"Core Event Vector Table" on page 4-40
0xFFE0 2028	EVT10 (IVG10)	"Core Event Vector Table" on page 4-40
0xFFE0 202C	EVT11 (IVG11)	"Core Event Vector Table" on page 4-40
0xFFE0 2030	EVT12 (IVG12)	"Core Event Vector Table" on page 4-40

Table A-3. Interrupt Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 2034	EVT13 (IVG13)	"Core Event Vector Table" on page 4-40
0xFFE0 2038	EVT14 (IVG14)	"Core Event Vector Table" on page 4-40
0xFFE0 203C	EVT15 (IVG15)	"Core Event Vector Table" on page 4-40
0xFFE0 2104	IMASK	"Core Interrupt Mask Register (IMASK)" on page 4-36
0xFFE0 2108	IPEND	"Core Interrupt Pending Register (IPEND)" on page 4-38
0xFFE0 2110	IPRIO	"Interrupt Priority Register (IPRIO) and Write Buffer Depth" on page 6-37
0xFFE0 210C	ILAT	"Core Interrupt Latch Register (ILAT)" on page 4-37

Core Timer Registers

Core Timer registers (0xFFE0 3000 – 0xFFE0 300C) are shown in Table A-4.

Table A-4. Core Timer Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 3000	TCNTL	"Core Timer Control Register (TCNTL)" on page 15-50
0xFFE0 3004	TPERIOD	"Core Timer Period Register (TPERIOD)" on page 15-52
0xFFE0 3008	TSCALE	"Core Timer Scale Register (TSCALE)" on page 15-53
0xFFE0 300C	TCOUNT	"Core Timer Count Register (TCOUNT)" on page 15-51

Debug, MP, and Emulation Unit Registers

Debug, MP, and Emulation Unit registers (0xFFE0 5000 – 0xFFE0 5008) are shown in Table A-5.

Table A-5. Debug and Emulation Unit Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 5000	DSPID	"DSP Device ID Register (DSPID)" on page 18-26

Trace Unit Registers

Trace Unit registers (0xFFE0 6000 – 0xFFE0 6100) are shown in Table A-6.

Table A-6. Trace Unit Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 6000	TBUFCTL	"Trace Buffer Control Register (TBUFCTL)" on page 18-16
0xFFE0 6004	TBUFSTAT	"Trace Buffer Status Register (TBUFSTAT)" on page 18-17
0xFFE0 6100	TBUF	"Trace Buffer Register (TBUF)" on page 18-17

Watchpoint and Patch Registers

Watchpoint and Patch registers (0xFFE0 7000 – 0xFFE0 7200) are shown in Table A-7.

Table A-7. Watchpoint and Patch Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 7000	WPIACTL	"Instruction Watchpoint Address Control Register (WPIACTL)" on page 18-7
0xFFE0 7040	WPIA0	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5
0xFFE0 7044	WPIA1	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5
0xFFE0 7048	WPIA2	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5
0xFFE0 704C	WPIA3	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5

Table A-7. Watchpoint and Patch Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFE0 7050	WPIA4	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5
0xFFE0 7054	WPIA5	"Instruction Watchpoint Address Registers (WPIAn)" on page 18-5
0xFFE0 7080	WPIACNT0	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 7084	WPIACNT1	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 7088	WPIACNT2	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 708C	WPIACNT3	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 7090	WPIACNT4	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 7094	WPIACNT5	"Instruction Watchpoint Address Count Registers (WPIACNTn)" on page 18-6
0xFFE0 7100	WPDACTL	"Data Watchpoint Address Control Register (WPDACTL)" on page 18-12
0xFFE0 7140	WPDA0	"Data Watchpoint Address Registers (WPDAn)" on page 18-11
0xFFE0 7144	WPDA1	"Data Watchpoint Address Registers (WPDAn)" on page 18-11
0xFFE0 7180	WPDACNT0	"Data Watchpoint Address Count Value Registers (WPDACNTn)" on page 18-11
0xFFE0 7184	WPDACNT1	"Data Watchpoint Address Count Value Registers (WPDACNTn)" on page 18-11
0xFFE0 7200	WPSTAT	"Watchpoint Status Register (WPSTAT)" on page 18-12

Performance Monitor Registers

Performance Monitor registers (0xFFE0 8000 – 0xFFE0 8104) are shown in Table A-8.

Table A-8. Performance Monitor Registers

Memory-Mapped Address	Register Name	See Section
0xFFE0 8000	PFCTL	"Performance Monitor Control Register (PFCTL)" on page 18-20
0xFFE0 8100	PFCNTR0	"Performance Monitor Counter Registers (PFCNTRn)" on page 18-19
0xFFE0 8104	PFCNTR1	"Performance Monitor Counter Registers (PFCNTRn)" on page 18-19

B SYSTEM MMR ASSIGNMENTS

These notes provide general information about the system memory-mapped registers (MMRs):

- The system MMR address range is 0xFFC0 0000–0xFFDF FFFF.
- All system MMRs are either 16 bits or 32 bits wide. MMRs that are 16 bits wide must be accessed with 16-bit read or write operations. MMRs that are 32 bits wide must be accessed with 32-bit read or write operations. Check the description of the MMR to determine whether a 16-bit or a 32-bit access is required.
- All system MMR space that is not defined in this appendix is reserved for internal use only.

This appendix lists MMR addresses and register names. To find more information about an MMR, refer to the page shown in the “See Section” column. When viewing the PDF version of this document, click a reference in the “See Section” column to jump to the additional information about the MMR.

Dynamic Power Management Registers

The Dynamic Power Management registers (address 0xFFC0 0000 to 0xFFC0 00FF) are described in Table B-1.

Table B-1. Dynamic Power Management Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0000	PLL_CTL	"PLL Control Register (PLL_CTL)" on page 8-7
0xFFC0 0004	PLL_DIV	"PLL Divide Register (PLL_DIV)" on page 8-6
0xFFC0 0008	VR_CTL	"Voltage Regulator Control Register (VR_CTL)" on page 8-24
0xFFC0 000C	PLL_STAT	"PLL Status Register (PLL_STAT)" on page 8-9
0xFFC0 0010	PLL_LOCKCNT	"PLL Lock Count Register (PLL_LOCKCNT)" on page 8-11

System Reset and Interrupt Control Registers

The System Interrupt Controller registers (address 0xFFC0 0100 to 0xFFC0 01FF, and 0xFFC0 1100 to 0xFFC0 11FF) are described in Table B-2.

Table B-2. System Interrupt Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0100	SICA_SWRST	"Software Reset Register (SICA_SWRST)" on page 3-17
0xFFC0 0104	SICA_SYSCR	"System Reset Configuration Register A (SICA_SYSCR)" on page 3-14
0xFFC0 1104	SICB_SYSCR	"System Reset Configuration Register B (SICB_SYSCR)" on page 3-15

Table B-2. System Interrupt Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 010C	SICA_IMASK0	"System Interrupt Mask Registers (SIC_IMASKx)" on page 4-29
0xFFC0 0110	SICA_IMASK1	"System Interrupt Mask Registers (SIC_IMASKx)" on page 4-29
0xFFC0 0124	SICA_IAR0	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0128	SICA_IAR1	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 012C	SICA_IAR2	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0130	SICA_IAR3	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0134	SICA_IAR4	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0138	SICA_IAR5	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 013C	SICA_IAR6	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0140	SICA_IAR7	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 0114	SICA_ISR0	"System Interrupt Status Registers (SIC_ISRx)" on page 4-27
0xFFC0 0118	SICA_ISR1	"System Interrupt Status Registers (SIC_ISRx)" on page 4-27
0xFFC0 011C	SICA_IWR0	"System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26
0xFFC0 0120	SICA_IWR1	"System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26
0xFFC0 110C	SICB_IMASK0	"System Interrupt Mask Registers (SIC_IMASKx)" on page 4-29
0xFFC0 1110	SICB_IMASK1	"System Interrupt Mask Registers (SIC_IMASKx)" on page 4-29

Table B-2. System Interrupt Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 1124	SICB_IAR0	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1128	SICB_IAR1	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 112C	SICB_IAR2	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1130	SICB_IAR3	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1134	SICB_IAR4	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1138	SICB_IAR5	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 113C	SICB_IAR6	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1140	SICB_IAR7	"System Interrupt Assignment Registers (SIC_IARx)" on page 4-30
0xFFC0 1114	SICB_ISR0	"System Interrupt Status Registers (SIC_ISRx)" on page 4-27
0xFFC0 1118	SICB_ISR1	"System Interrupt Status Registers (SIC_ISRx)" on page 4-27
0xFFC0 111C	SICB_IWR0	"System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26
0xFFC0 1120	SICB_IWR1	"System Interrupt Wakeup-Enable Registers (SIC_IWRx)" on page 4-26

Watchdog Timer Registers

The Watchdog Timer registers (address 0xFFC0 0200 to 0xFFC0 02FF and 0xFFC0 1200 to 0xFFC0 12FF) are described in Table B-3.

Table B-3. Watchdog Timer Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0200	WDOGA_CTL	"Watchdog Control Register (WDOG_CTL)" on page 15-56
0xFFC0 0204	WDOGA_CNT	"Watchdog Count Register (WDOG_CNT)" on page 15-54
0xFFC0 0208	WDOGA_STAT	"Watchdog Status Register (WDOG_STAT)" on page 15-55
0xFFC0 1200	WDOGB_CTL	"Watchdog Control Register (WDOG_CTL)" on page 15-56
0xFFC0 1204	WDOGB_CNT	"Watchdog Count Register (WDOG_CNT)" on page 15-54
0xFFC0 1208	WDOGB_STAT	"Watchdog Status Register (WDOG_STAT)" on page 15-55

Parallel Peripheral Interface (PPI) Registers

The Parallel Peripheral Interface (PPI) registers (address 0xFFC0 1000 to 0xFFC0 10FF and 0xFFC0 1300 to 0xFFC0 13FF) are described in Table B-4.

Table B-4. Parallel Peripheral Interface (PPI) Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 1000	PPI0_CONTROL	"PPI Control Register (PPIx_CONTROL)" on page 11-3
0xFFC0 1004	PPI0_STATUS	"PPI Status Register (PPIx_STATUS)" on page 11-9
0xFFC0 1008	PPI0_COUNT	"Transfer Count Register (PPIx_COUNT)" on page 11-12

Table B-4. Parallel Peripheral Interface (PPI) Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 100C	PPI0_DELAY	"Delay Count Register (PPIx_DELAY)" on page 11-11
0xFFC0 1010	PPI0_FRAME	"Lines Per Frame Register (PPIx_FRAME)" on page 11-13
0xFFC0 1300	PPI1_CONTROL	"PPI Control Register (PPIx_CONTROL)" on page 11-3
0xFFC0 1304	PPI1_STATUS	"PPI Status Register (PPIx_STATUS)" on page 11-9
0xFFC0 1308	PPI1_COUNT	"Transfer Count Register (PPIx_COUNT)" on page 11-12
0xFFC0 130C	PPI1_DELAY	"Delay Count Register (PPIx_DELAY)" on page 11-11
0xFFC0 1310	PPI1_FRAME	"Lines Per Frame Register (PPIx_FRAME)" on page 11-13

UART Controller Registers

The UART Controller registers (address 0xFFC0 0400 to 0xFFC0 04FF) are described in Table B-5.

Table B-5. UART Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0400	UART_THR	"UART Transmit Holding Register (UART_THR)" on page 13-7
0xFFC0 0400	UART_RBR	"UART Receive Buffer Register (UART_RBR)" on page 13-8
0xFFC0 0400	UART_DLL	"UART Divisor Latch Registers (UART_DLL, UART_DLH)" on page 13-12
0xFFC0 0404	UART_DLH	"UART Divisor Latch Registers (UART_DLL, UART_DLH)" on page 13-12
0xFFC0 0404	UART_IER	"UART Interrupt Enable Register (UART_IER)" on page 13-9
0xFFC0 0408	UART_IIR	"UART Interrupt Identification Register (UART_IIR)" on page 13-10

Table B-5. UART Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 040C	UART_LCR	"UART Line Control Register (UART_LCR)" on page 13-4
0xFFC0 0410	UART_MCR	"Modem Control Register (UART_MCR)" on page 13-5
0xFFC0 0414	UART_LSR	"UART Line Status Register (UART_LSR)" on page 13-6
0xFFC0 041C	UART_SCR	"UART Scratch Register (UART_SCR)" on page 13-14
0xFFC0 0424	UART_GCTL	"UART Global Control Register (UART_GCTL)" on page 13-14

SPI Controller Registers

The SPI Controller registers (address `0xFFC0 0500` to `0xFFC0 05FF`) are described in Table B-6.

Table B-6. SPI Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0500	SPI_CTL	"SPI Control Register (SPI_CTL)" on page 10-8
0xFFC0 0504	SPI_FLG	"SPI Flag Register (SPI_FLG)" on page 10-11
0xFFC0 0508	SPI_STAT	"SPI Status Register (SPI_STAT)" on page 10-15
0xFFC0 050C	SPI_TDBR	"SPI Transmit Data Buffer Register (SPI_TDBR)" on page 10-17
0xFFC0 0510	SPI_RDBR	"SPI Receive Data Buffer Register (SPI_RDBR)" on page 10-18
0xFFC0 0514	SPI_BAUD	"SPI Baud Rate Register (SPI_BAUD)" on page 10-7
0xFFC0 0518	SPI_SHADOW	"SPI RDBR Shadow Register (SPI_SHADOW)" on page 10-18

Timer Registers

The Timer registers are described in Table B-7 (address 0xFFC0 0600 to 0xFFC0 06FF and 0xFFC0 1600 to 0xFFC0 16FF).

Table B-7. Timer Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0600	TIMER0_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0604	TIMER0_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0608	TIMER0_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 060C	TIMER0_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0610	TIMER1_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0614	TIMER1_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0618	TIMER1_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 061C	TIMER1_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0620	TIMER2_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0624	TIMER2_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0628	TIMER2_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15

Table B-7. Timer Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 062C	TIMER2_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0630	TIMER3_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0634	TIMER3_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0638	TIMER3_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 063C	TIMER3_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0640	TIMER4_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0644	TIMER4_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0648	TIMER4_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 064C	TIMER4_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0650	TIMER5_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0654	TIMER5_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0658	TIMER5_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15

Table B-7. Timer Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 065C	TIMER5_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0660	TIMER6_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0664	TIMER6_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0668	TIMER6_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 066C	TIMER6_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0670	TIMER7_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 0674	TIMER7_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 0678	TIMER7_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 067C	TIMER7_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 1600	TIMER8_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 1604	TIMER8_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 1608	TIMER8_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15

Table B-7. Timer Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 160C	TIMER8_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 1610	TIMER9_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 1614	TIMER9_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 1618	TIMER9_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 161C	TIMER9_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 1620	TIMER10_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 1624	TIMER10_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 1628	TIMER10_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 162C	TIMER10_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 1630	TIMER11_CONFIG	"Timer Configuration (TIMERx_CONFIG) Registers" on page 15-12
0xFFC0 1634	TIMER11_COUNTER	"Timer Counter (TIMERx_COUNTER) Registers" on page 15-14
0xFFC0 1638	TIMER11_PERIOD	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15

Table B-7. Timer Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 163C	TIMER11_WIDTH	"Timer Period (TIMERx_PERIOD) and Timer Pulse Width (TIMERx_WIDTH) Registers" on page 15-15
0xFFC0 0680	TMRS8_ENABLE	"Timer Enable (TMRSy_ENABLE) Registers" on page 15-5
0xFFC0 0684	TMRS8_DISABLE	"Timer Disable (TMRSy_DISABLE) Registers" on page 15-5
0xFFC0 0688	TMRS8_STATUS	"Timer Status (TMRSy_STATUS) Registers" on page 15-8
0xFFC0 1640	TMRS4_ENABLE	"Timer Enable (TMRSy_ENABLE) Registers" on page 15-5
0xFFC0 1644	TMRS4_DISABLE	"Timer Disable (TMRSy_DISABLE) Registers" on page 15-5
0xFFC0 1648	TMRS4_STATUS	"Timer Status (TMRSy_STATUS) Registers" on page 15-8

Programmable Flag Registers

The Programmable Flags registers (address 0xFFC0 0700 to 0xFFC0 07FF and 0xFFC0 1500 to 0xFFC0 15FF and 0xFFC0 1700 to 0xFFC0 17FF) are described in Table B-8.

Table B-8. Programmable Flags Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0700	FIO0_FLAG_D	"Flag Data Register (FIOn_FLAG_D)" on page 14-12
0xFFC0 0704	FIO0_FLAG_C	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13

Table B-8. Programmable Flags Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 0708	FIO0_FLAG_S	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 070C	FIO0_FLAG_T	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 0710	FIO0_MASKA_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0714	FIO0_MASKA_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0718	FIO0_MASKA_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 071C	FIO0_MASKA_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0720	FIO0_MASKB_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0724	FIO0_MASKB_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0728	FIO0_MASKB_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19

Table B-8. Programmable Flags Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 072C	FIO0_MASKB_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 0730	FIO0_DIR	"Flag Direction Register (FIOn_DIR)" on page 14-8
0xFFC0 0734	FIO0_POLAR	"Flag Polarity Register (FIOn_POLAR)" on page 14-33
0xFFC0 0738	FIO0_EDGE	"Flag Interrupt Sensitivity Register (FIOn_EDGE)" on page 14-35
0xFFC0 073C	FIO0_BOTH	"Flag Set on Both Edges Register (FIOn_BOTH)" on page 14-37
0xFFC0 0740	FIO0_INEN	"Flag Input Enable Register (FIOn_INEN)" on page 14-39
0xFFC0 1500	FIO1_FLAG_D	"Flag Data Register (FIOn_FLAG_D)" on page 14-12
0xFFC0 1504	FIO1_FLAG_C	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 1508	FIO1_FLAG_S	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 150C	FIO1_FLAG_T	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 1510	FIO1_MASKA_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1514	FIO1_MASKA_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19

Table B-8. Programmable Flags Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 1518	FIO1_MASKA_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 151C	FIO1_MASKA_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1520	FIO1_MASKB_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1524	FIO1_MASKB_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1528	FIO1_MASKB_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 152C	FIO1_MASKB_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1530	FIO1_DIR	"Flag Direction Register (FIOn_DIR)" on page 14-8
0xFFC0 1534	FIO1_POLAR	"Flag Polarity Register (FIOn_POLAR)" on page 14-33
0xFFC0 1538	FIO1_EDGE	"Flag Interrupt Sensitivity Register (FIOn_EDGE)" on page 14-35
0xFFC0 153C	FIO1_BOTH	"Flag Set on Both Edges Register (FIOn_BOTH)" on page 14-37
0xFFC0 1540	FIO1_INEN	"Flag Input Enable Register (FIOn_INEN)" on page 14-39
0xFFC0 1700	FIO2_FLAG_D	"Flag Data Register (FIOn_FLAG_D)" on page 14-12

Table B-8. Programmable Flags Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 1704	FIO2_FLAG_C	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 1708	FIO2_FLAG_S	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 170C	FIO2_FLAG_T	"Flag Set (FIOn_FLAG_S), Flag Clear (FIOn_FLAG_C), and Flag Toggle (FIOn_FLAG_T) Registers" on page 14-13
0xFFC0 1710	FIO2_MASKA_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1714	FIO2_MASKA_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1718	FIO2_MASKA_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 171C	FIO2_MASKA_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1720	FIO2_MASKB_D	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1724	FIO2_MASKB_C	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19

Table B-8. Programmable Flags Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 1728	FIO2_MASKB_S	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 172C	FIO2_MASKB_T	"Flag Mask Interrupt Registers (FIOn_MASKA_D, FIOn_MASKA_C, FIOn_MASKA_S,FIOn_MASKA_T, FIOn_MASKB_D, FIOn_MASKB_C, FIOn_MASKB_S, FIOn_MASKB_T)" on page 14-19
0xFFC0 1730	FIO2_DIR	"Flag Direction Register (FIOn_DIR)" on page 14-8
0xFFC0 1734	FIO2_POLAR	"Flag Polarity Register (FIOn_POLAR)" on page 14-33
0xFFC0 1738	FIO2_EDGE	"Flag Interrupt Sensitivity Register (FIOn_EDGE)" on page 14-35
0xFFC0 173C	FIO2_BOTH	"Flag Set on Both Edges Register (FIOn_BOTH)" on page 14-37
0xFFC0 1740	FIO2_INEN	"Flag Input Enable Register (FIOn_INEN)" on page 14-39

SPORT0 Controller Registers

The SPORT0 Controller registers (address `0xFFC0 0800` to `0xFFC0 08FF`) are described in Table B-9.

Table B-9. SPORT0 Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0800	SPORT0_TCR1	"SPORTx_TCR1 and SPORTx_TCR2 Registers" on page 12-11
0xFFC0 0804	SPORT0_TCR2	"SPORTx_TCR1 and SPORTx_TCR2 Registers" on page 12-11

Table B-9. SPORT0 Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 0808	SPORT0_TCLKDIV	"SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers" on page 12-27
0xFFC0 080C	SPORT0_TFSDIV	"SPORTx_TFSDIV and SPORTx_RFSDIV Register" on page 12-28
0xFFC0 0810	SPORT0_TX	"SPORTx_TX Register" on page 12-21
0xFFC0 0818	SPORT0_RX	"SPORTx_RX Register" on page 12-23
0xFFC0 0820	SPORT0_RCR1	"SPORTx_RCR1 and SPORTx_RCR2 Registers" on page 12-16
0xFFC0 0824	SPORT0_RCR2	"SPORTx_RCR1 and SPORTx_RCR2 Registers" on page 12-16
0xFFC0 0828	SPORT0_RCLKDIV	"SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers" on page 12-27
0xFFC0 082C	SPORT0_RFSDIV	"SPORTx_TFSDIV and SPORTx_RFSDIV Register" on page 12-28
0xFFC0 0830	SPORT0_STAT	"SPORTx_STAT Register" on page 12-25
0xFFC0 0834	SPORT0_CHNL	"SPORTx_CHNL Register" on page 12-54
0xFFC0 0838	SPORT0_MCMC1	"SPORTx_MCMCn Registers" on page 12-49
0xFFC0 083C	SPORT0_MCMC2	"SPORTx_MCMCn Registers" on page 12-49
0xFFC0 0840	SPORT0_MTCS0	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 0844	SPORT0_MTCS1	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 0848	SPORT0_MTCS2	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 084C	SPORT0_MTCS3	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 0850	SPORT0_MRCS0	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 0854	SPORT0_MRCS1	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 0858	SPORT0_MRCS2	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 085C	SPORT0_MRCS3	"SPORTx_MRCSn Registers" on page 12-57

SPORT1 Controller Registers

The SPORT1 Controller registers (address 0xFFC0 0900 to 0xFFC0 09FF) are described in Table B-10.

Table B-10. SPORT 1 Controller Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0900	SPORT1_TCR1	"SPORTx_TCR1 and SPORTx_TCR2 Registers" on page 12-11
0xFFC0 0904	SPORT1_TCR2	"SPORTx_TCR1 and SPORTx_TCR2 Registers" on page 12-11
0xFFC0 0908	SPORT1_TCLKDIV	"SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers" on page 12-27
0xFFC0 090C	SPORT1_TFSDIV	"SPORTx_TFSDIV and SPORTx_RFSDIV Register" on page 12-28
0xFFC0 0910	SPORT1_TX	"SPORTx_TX Register" on page 12-21
0xFFC0 0918	SPORT1_RX	"SPORTx_RX Register" on page 12-23
0xFFC0 0920	SPORT1_RCR1	"SPORTx_RCR1 and SPORTx_RCR2 Registers" on page 12-16
0xFFC0 0924	SPORT1_RCR2	"SPORTx_RCR1 and SPORTx_RCR2 Registers" on page 12-16
0xFFC0 0928	SPORT1_RCLKDIV	"SPORTx_TCLKDIV and SPORTx_RCLKDIV Registers" on page 12-27
0xFFC0 092C	SPORT1_RFSDIV	"SPORTx_TFSDIV and SPORTx_RFSDIV Register" on page 12-28
0xFFC0 0930	SPORT1_STAT	"SPORTx_STAT Register" on page 12-25
0xFFC0 0934	SPORT1_CHNL	"SPORTx_CHNL Register" on page 12-54
0xFFC0 0938	SPORT1_MCMC1	"SPORTx_MCMCn Registers" on page 12-49
0xFFC0 093C	SPORT1_MCMC2	"SPORTx_MCMCn Registers" on page 12-49
0xFFC0 0940	SPORT1_MTCS0	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 0944	SPORT1_MTCS1	"SPORTx_MTCSn Registers" on page 12-58

Table B-10. SPORT 1 Controller Registers (Cont'd)

Memory-Mapped Address	Register Name	See Section
0xFFC0 0948	SPORT1_MTCS2	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 094C	SPORT1_MTCS3	"SPORTx_MTCSn Registers" on page 12-58
0xFFC0 0950	SPORT1_MRCS0	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 0954	SPORT1_MRCS1	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 0958	SPORT1_MRCS2	"SPORTx_MRCSn Registers" on page 12-57
0xFFC0 095C	SPORT1_MRCS3	"SPORTx_MRCSn Registers" on page 12-57

DMA Control Registers

DMA1/MDMA1 Control Registers	0xFFC0 1C00 - 0xFFC0 1FFF
DMA2/MDMA2 Control Registers	0xFFC0 0C00 - 0xFFC0 0FFF
IMDMA Control Registers	0xFFC0 1800 - 0xFFC0 19FF

Since each DMA channel has an identical MMR set, with fixed offsets from the base address associated with that DMA channel, it is convenient to view the MMR information as provided in Table B-11 and Table B-12. Table B-11 identifies the "Base Address" of each DMA channel within each DMA Controller, as well as the register prefix that identifies that particular channel. Table B-12 then provides the register suffix and provides its offset from the Base Address.

As an example, the DMA Channel 0 `Y_MODIFY` register in DMA Controller 1 is called `DMA1_0_Y_MODIFY`, and its address is `0xFFC0 1C1C`. Likewise, the Memory DMA Stream 0 Source Current Address register in DMA Controller 2 is called `MDMA2_S0_CURR_ADDR`, and its address is `0xFFC0 0F64`.

Table B-11. DMA Channel Base Address

DMA Controller	DMA Channel Identifier	MMR Base Address	Register Prefix
1	0	0xFFC0 1C00	DMA1_0_
1	1	0xFFC0 1C40	DMA1_1_
1	2	0xFFC0 1C80	DMA1_2_
1	3	0xFFC0 1CC0	DMA1_3_
1	4	0xFFC0 1D00	DMA1_4_
1	5	0xFFC0 1D40	DMA1_5_
1	6	0xFFC0 1D80	DMA1_6_
1	7	0xFFC0 1DC0	DMA1_7_
1	8	0xFFC0 1E00	DMA1_8_
1	9	0xFFC0 1E40	DMA1_9_
1	10	0xFFC0 1E80	DMA1_10_
1	11	0xFFC0 1EC0	DMA1_11_
1	Mem DMA Stream 0 Destination	0xFFC0 1F00	MDMA1_D0_
1	Mem DMA Stream 0 Source	0xFFC0 1F40	MDMA1_S0_
1	Mem DMA Stream 1 Destination	0xFFC0 1F80	MDMA1_D1_
1	Mem DMA Stream 1 Source	0xFFC0 1FC0	MDMA1_S1_
2	0	0xFFC0 0C00	DMA2_0_
2	1	0xFFC0 0C40	DMA2_1_
2	2	0xFFC0 0C80	DMA2_2_
2	3	0xFFC0 0CC0	DMA2_3_
2	4	0xFFC0 0D00	DMA2_4_
2	5	0xFFC0 0D40	DMA2_5_
2	6	0xFFC0 0D80	DMA2_6_
2	7	0xFFC0 0DC0	DMA2_7_
2	8	0xFFC0 0E00	DMA2_8_
2	9	0xFFC0 0E40	DMA2_9_

Table B-11. DMA Channel Base Address (Cont'd)

DMA Controller	DMA Channel Identifier	MMR Base Address	Register Prefix
2	10	0xFFC0 0E80	DMA2_10_
2	11	0xFFC0 0EC0	DMA2_11_
2	Mem DMA Stream 0 Destination	0xFFC0 0F00	MDMA2_D0_
2	Mem DMA Stream 0 Source	0xFFC0 0F40	MDMA2_S0_
2	Mem DMA Stream 1 Destination	0xFFC0 0F80	MDMA2_D1_
2	Mem DMA Stream 1 Source	0xFFC0 0FC0	MDMA2_S1_
IMDMA	IMDMA Stream 0 Destination	0xFFC0 1800	IMDMA_D0
IMDMA	IMDMA Stream 0 Source	0xFFC0 1840	IMDMA_S0
IMDMA	IMDMA Stream 1 Destination	0xFFC0 1880	IMDMA_D1
IMDMA	IMDMA Stream 1 Source	0xFFC0 18C0	IMDMA_S1

Table B-12. DMA Channel Offset

Register Suffix	Offset from Base	See Section
NEXT_DESC_PTR	0x00	"Next Descriptor Pointer Register (DMAx_y_NEXT_DESC_PTR/ MDMA_yy_NEXT_DESC_PTR)" on page 9-10
START_ADDR	0x04	"Start Address Register (DMAx_y_START_ADDR/ MDMAx_yy_START_ADDR)" on page 9-11
CONFIG	0x08	"DMA Configuration Register (DMAx_y_CONFIG/MDMAx_yy_CONFIG)" on page 9-12
X_COUNT	0x10	"Inner Loop Count Register (DMAx_y_X_COUNT/MDMAx_yy_X_COUNT)" on page 9-16
X_MODIFY	0x14	"Inner Loop Address Increment Register (DMAx_y_X_MODIFY/MDMAx_yy_X_MODIFY)" on page 9-16
Y_COUNT	0x18	"Outer Loop Count Register (DMAx_y_Y_COUNT/MDMAx_yy_Y_COUNT)" on page 9-17

Table B-12. DMA Channel Offset (Cont'd)

Register Suffix	Offset from Base	See Section
Y_MODIFY	0x1C	"Outer Loop Address Increment Register (DMAx_y_Y_MODIFY/MDMAx_yy_Y_MODIFY)" on page 9-18
CURR_DESC_PTR	0x20	"Current Descriptor Pointer Register (DMAx_y_CURR_DESC_PTR/ MDMAx_yy_CURR_DESC_PTR)" on page 9-19
CURR_ADDR	0x24	"Current Address Register (DMAx_y_CURR_ADDR/ MDMAx_yy_CURR_ADDR)" on page 9-20
IRQ_STATUS	0x28	"Interrupt Status Register (DMAx_y_IRQ_STATUS/MDMAx_yy_IRQ_STATUS)" on page 9-26
PERIPHERAL_MAP	0x2C	"Peripheral Map Register (DMAx_y_PERIPHERAL_MAP/ MDMAx_yy_PERIPHERAL_MAP)" on page 9-23
CURR_X_COUNT	0x30	"Current Inner Loop Count Register (DMAx_y_CURR_X_COUNT/ MDMAx_yy_CURR_X_COUNT)" on page 9-21
CURR_Y_COUNT	0x38	"Current Outer Loop Count Register (DMAx_y_CURR_Y_COUNT/ MDMAx_yy_CURR_Y_COUNT)" on page 9-22

DMA Traffic Control Registers

The DMA Traffic Control registers are described in Table B-13.

Table B-13. DMA Traffic Control Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 1B0C	DMA1_TC_PER	Figure 9-17 on page 9-49
0xFFC0 1B10	DMA1_TC_CNT	Figure 9-18 on page 9-49
0xFFC0 0B0C	DMA2_TC_PER	Figure 9-17 on page 9-49
0xFFC0 0B10	DMA2_TC_CNT	Figure 9-18 on page 9-49

External Bus Interface Unit Registers

The External Bus Interface Unit registers (address 0xFFC0 0A00 to 0xFFC0 0aFF) are described in Table B-14.

Table B-14. External Bus Interface Unit Registers

Memory-Mapped Address	Register Name	See Section
0xFFC0 0A00	EBIU_AMGCTL	"Asynchronous Memory Global Control Register (EBIU_AMGCTL)" on page 16-9
0xFFC0 0A04	EBIU_AMBCTL0	"Asynchronous Memory Bank Control Registers (EBIU_AMBCTL0, EBIU_AMBCTL1)" on page 16-12
0xFFC0 0A08	EBIU_AMBCTL1	"Asynchronous Memory Bank Control Registers (EBIU_AMBCTL0, EBIU_AMBCTL1)" on page 16-12
0xFFC0 0A10	EBIU_SDGCTL	"SDRAM Memory Global Control Register (EBIU_SDGCTL)" on page 16-34
0xFFC0 0A14	EBIU_SDBCTL	"SDRAM Memory Bank Control Register (EBIU_SDBCTL)" on page 16-46
0xFFC0 0A18	EBIU_SDRRC	"SDRAM Refresh Rate Control Register (EBIU_SDRRC)" on page 16-50
0xFFC0 0A1C	EBIU_SDSTAT	"SDRAM Control Status Register (EBIU_SDSTAT)" on page 16-49

C NUMERIC FORMATS

Blackfin family processors support 8-, 16-, 32-, and 40-bit fixed-point data in hardware. Special features in the computation units allow support of other formats in software. This appendix describes various aspects of these data formats. It also describes how to implement a block floating-point format in software.

Unsigned or Signed: Two's-Complement Format

Unsigned integer numbers are positive, and no sign information is contained in the bits. Therefore, the value of an unsigned integer is interpreted in the usual binary sense. The least significant words of multiple-precision numbers are treated as unsigned numbers.

Signed numbers supported by the Blackfin family are in two's-complement format. Signed-magnitude, one's-complement, binary-coded decimal (BCD) or excess-n formats are not supported.

Integer or Fractional

The Blackfin family supports both fractional and integer data formats. In an integer, the radix point is assumed to lie to the right of the least significant bit (LSB), so that all magnitude bits have a weight of 1 or greater. This format is shown in Figure C-1. Note in two's-complement format, the sign bit has a negative weight.

In a fractional format, the assumed radix point lies within the number, so that some or all of the magnitude bits have a weight of less than 1. In the format shown in Figure C-2, the assumed radix point lies to the left of the three LSBs, and the bits have the weights indicated.

The native formats for the Blackfin processor family are a signed fractional 1.M format and an unsigned fractional 0.N format, where N is the number of bits in the data word and M = N – 1.

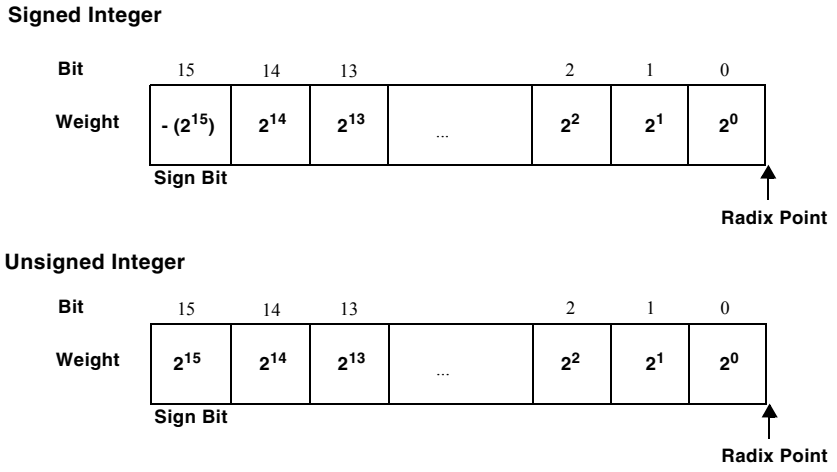

Figure C-1. Integer Format

The notation used to describe a format consists of two numbers separated by a period (.); the first number is the number of bits to the left of the radix point, while the second is the number of bits to the right of the radix point. For example, 16.0 format is an integer format; all bits lie to the left of the radix point. The format in Figure C-2 is 13.3.

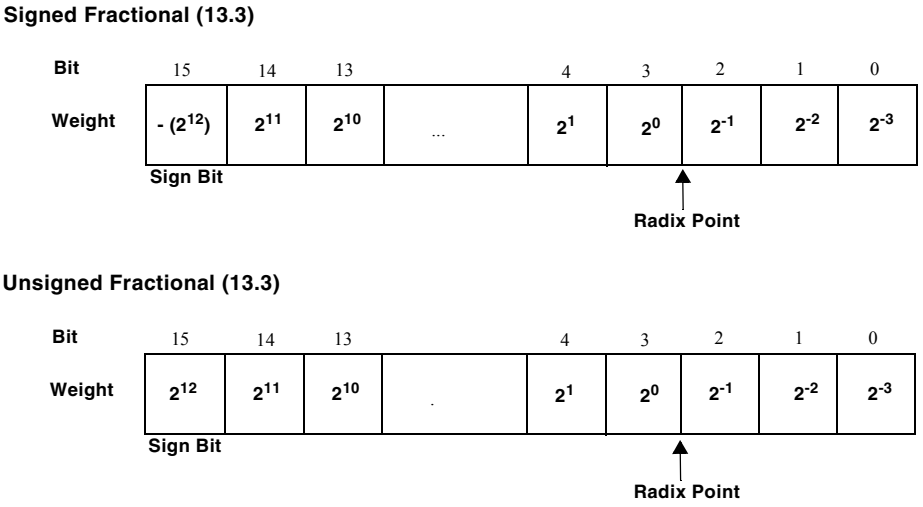

Figure C-2. Example of Fractional Format

Table C-1 shows the ranges of signed numbers representable in the fractional formats that are possible with 16 bits.

Table C-1. Fractional Formats and Their Ranges

Format	# of Integer Bits	# of Fractional Bits	Max Positive Value (0x7FFF) In Decimal	Max Negative Value (0x8000) In Decimal	Value of 1 LSB (0x0001) In Decimal
1.15	1	15	0.999969482421875	–1.0	0.000030517578125
2.14	2	14	1.999938964843750	–2.0	0.000061035156250
3.13	3	13	3.999877929687500	–4.0	0.000122070312500
4.12	4	12	7.999755859375000	–8.0	0.000244140625000
5.11	5	11	15.999511718750000	–16.0	0.000488281250000
6.10	6	10	31.999023437500000	–32.0	0.000976562500000
7.9	7	9	63.998046875000000	–64.0	0.001953125000000
8.8	8	8	127.996093750000000	–128.0	0.003906250000000
9.7	9	7	255.992187500000000	–256.0	0.007812500000000
10.6	10	6	511.984375000000000	–512.0	0.015625000000000
11.5	11	5	1023.968750000000000	–1024.0	0.031250000000000
12.4	12	4	2047.937500000000000	–2048.0	0.062500000000000
13.3	13	3	4095.875000000000000	–4096.0	0.125000000000000
14.2	14	2	8191.750000000000000	–8192.0	0.250000000000000
15.1	15	1	16383.500000000000000	–16384.0	0.500000000000000
16.0	16	0	32767.000000000000000	–32768.0	1.000000000000000

Binary Multiplication

In addition and subtraction, both operands must be in the same format (signed or unsigned, radix point in the same location), and the result format is the same as the input format. Addition and subtraction are performed the same way whether the inputs are signed or unsigned.

In multiplication, however, the inputs can have different formats, and the result depends on their formats. The Blackfin family assembly language allows you to specify whether the inputs are both signed, both unsigned, or one of each (mixed-mode). The location of the radix point in the result can be derived from its location in each of the inputs. This is shown in Figure C-3. The product of two 16-bit numbers is a 32-bit number. If the inputs' formats are M.N and P.Q, the product has the format (M + P).(N + Q). For example, the product of two 13.3 numbers is a 26.6 number. The product of two 1.15 numbers is a 2.30 number.

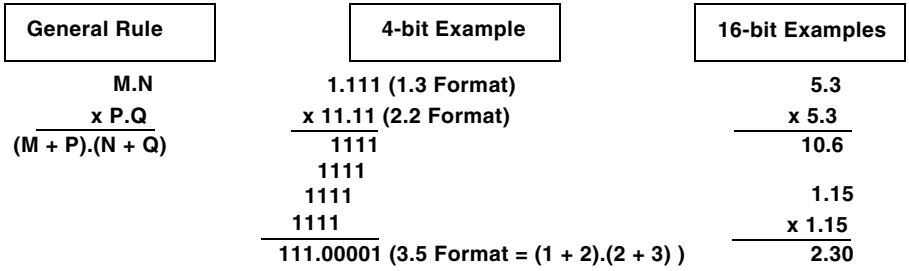

Figure C-3. Format of Multiplier Result

Fractional Mode And Integer Mode

A product of 2 two's-complement numbers has two sign bits. Since one of these bits is redundant, you can shift the entire result left one bit. Additionally, if one of the inputs was a 1.15 number, the left shift causes the result to have the same format as the other input (with 16 bits of additional precision). For example, multiplying a 1.15 number by a 5.11 number yields a 6.26 number. When shifted left one bit, the result is a 5.27 number, or a 5.11 number plus 16 LSBs.

The Blackfin family provides a means (a signed fractional mode) by which the multiplier result is always shifted left one bit before being written to

the result register. This left shift eliminates the extra sign bit when both operands are signed, yielding a result that is correctly formatted.

When both operands are in 1.15 format, the result is 2.30 (30 fractional bits). A left shift causes the multiplier result to be 1.31 which can be rounded to 1.15. Thus, if you use a signed fractional data format, it is most convenient to use the 1.15 format.

For more information about data formats, see the data formats listed in Table 2-2 on page 2-11.

Block Floating-Point Format

A block floating-point format enables a fixed-point processor to gain some of the increased dynamic range of a floating-point format without the overhead needed to do floating-point arithmetic. However, some additional programming is required to maintain a block floating-point format.

A floating-point number has an exponent that indicates the position of the radix point in the actual value. In block floating-point format, a set (block) of data values share a common exponent. A block of fixed-point values can be converted to block floating-point format by shifting each value left by the same amount and storing the shift value as the block exponent.

Typically, block floating-point format allows you to shift out non-significant MSBs (most significant bits), increasing the precision available in each value. Block floating-point format can also be used to eliminate the possibility of a data value overflowing. See Figure C-4. Each of the three data samples shown has at least two non-significant, redundant sign bits.

Each data value can grow by these two bits (two orders of magnitude) before overflowing. These bits are called guard bits.

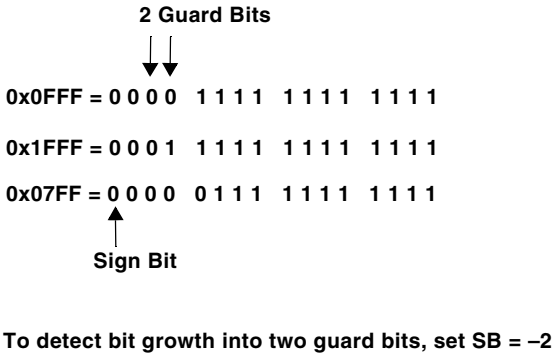

Figure C-4. Data With Guard Bits

If it is known that a process will not cause any value to grow by more than the two guard bits, then the process can be run without loss of data. Later, however, the block must be adjusted to replace the guard bits before the next process.

Figure C-5 shows the data after processing but before adjustment. The block floating-point adjustment is performed as follows.

- Assume the output of the `SIGNBITS` instruction is `SB` and `SB` is used as an argument in the `EXPADJ` instruction (see the *Blackfin Processor Instruction Set Reference* for the usage and syntax of these instructions). Initially, the value of `SB` is +2, corresponding to the two guard bits. During processing, each resulting data value is inspected by the `EXPADJ` instruction, which counts the number of redundant sign bits and adjusts `SB` if the number of redundant sign

bits is less than two. In this example, SB = +1 after processing, indicating the block of data must be shifted right one bit to maintain the two guard bits.

- If SB were 0 after processing, the block would have to be shifted two bits right. In either case, the block exponent is updated to reflect the shift.

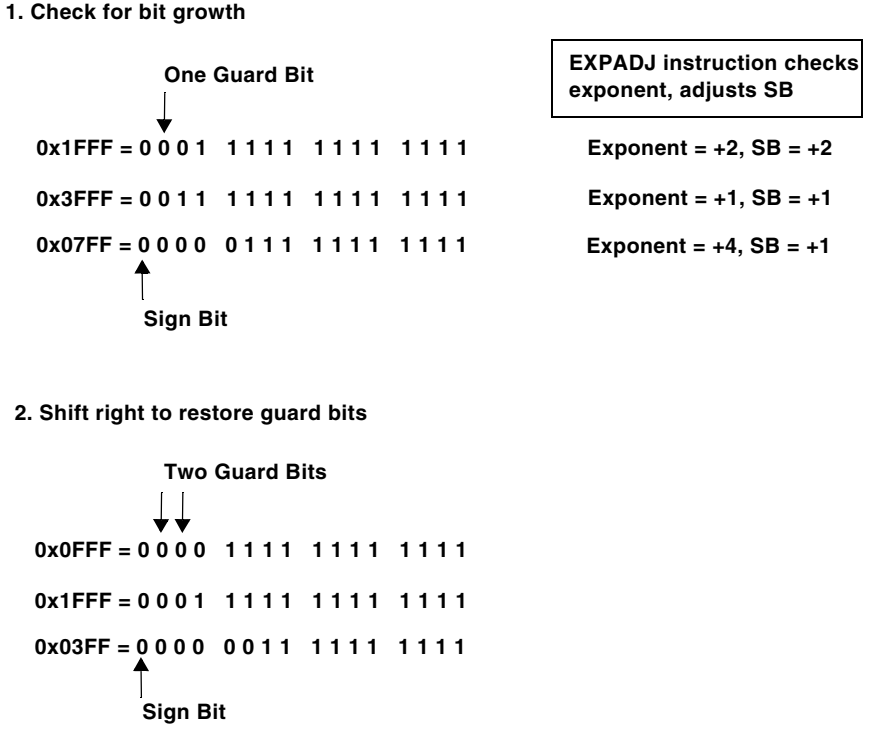

Figure C-5. Block Floating-point Adjustment

G GLOSSARY

ALU.

See *Arithmetic/Logic Unit*

AMC (Asynchronous Memory Controller).

A configurable memory controller for separate banks of memory, where each bank can be independently programmed with different timing parameters.

Arithmetic/Logic Unit (ALU).

A processor component that performs arithmetic, comparative, and logical functions.

asynchronous transfers.

Host accesses of the processor that occur intermittently rather than as a steady stream.

Bank Activate command.

The Bank Activate command causes the SDRAM to open an internal bank (specified by the bank address) in a row (specified by the row address). When the Bank Activate command is issued to the SDRAM, the SDRAM opens a new row address and the new bank address. The memory in the open internal bank and row is referred to as the open page. The Bank Activate command must be applied before a read or write command. The SDRAM controller (SDC) does not interleave SDRAM accesses, so only one internal bank in a row is open at a time.

base address.

The starting address of a circular buffer that the Data Address Generator (DAG) wraps around.

Base register.

A Data Address Generator (DAG) register that sets up the starting address for a circular buffer.

bit-reversed addressing.

The process in which the Data Address Generator (DAG) provides a bit-reversed address during a data move without reversing the stored address.

Boot memory space.

Internal memory space designated for a program that is loaded by the EPROM after powerup or after a software reset.

burst length.

The burst length determines the number of words that the SDRAM device stores or delivers after detecting a single write or read command, respectively. The burst length is selected by writing certain bits in the SDRAM's Mode register during the SDRAM powerup sequence.

The SDRAM controller (SDC) supports only Burst Length = 1 mode. During a burst to SDRAM, the SDC applies the read or write command every cycle and keeps accessing the data.

Burst Stop command.

The Burst Stop command is one of several ways to terminate a burst read or write operation.

Since the SDRAM burst length is always programmed to be 1, the SDRAM controller (SDC) does not support the Burst Stop command.

burst type.

The burst type determines the address order in which the SDRAM delivers burst data after detecting a read command, or stores burst data after detecting a write command. The burst type is programmed in the SDRAM during the SDRAM powerup sequence.

Since the SDRAM burst length is always programmed to be 1, the burst type does not matter. However, the SDRAM controller (SDC) always sets the burst type to sequential-accesses-only during the SDRAM powerup sequence.

cache block.

The smallest unit of memory that is transferred to/from the next level of memory from/to a cache as a result of a cache miss.

cache hit.

A memory access that is satisfied by a valid, present entry in the cache.

cache line.

Same as cache block. In this document, *cache line* is used for *cache block*.

cache miss.

A memory access that does not match any valid entry in the cache.

Cacheability Protection Lookaside Buffer (CPLB).

Storage area that describes the access characteristics of the core memory map.

CAM (Content Addressable Memory).

Also called Associative Memory. A memory device that includes comparison logic with each bit of storage. A datum is broadcast to all words in

memory; the datum is compared with the stored values; and values that match are flagged.

CAS (Column Address Strobe).

A signal sent from the Memory Management Unit (MMU) to a DRAM device to indicate that the column address lines are valid.

CAS latency (also t_{AA}, t_{CAC}, CL).

The Column Address Strobe (CAS) latency is the delay in clock cycles between when the SDRAM detects the read command and when it provides the data at its output pins. The CAS latency is programmed in the SDRAM Mode register during the powerup sequence.

The speed grade of the SDRAM and the `SCLK[0]` frequency determine the value of the CAS latency. The SDRAM controller (SDC) can support CAS latency of 2 or 3 clock cycles. The selected CAS latency value must be programmed into the SDRAM Memory Global Control register (`EBIU_SDGCTL`) before the SDRAM powerup sequence. See "SDRAM Memory Global Control Register (EBIU_SDGCTL)" on page 16-34.

CBR (CAS Before RAS) memory refresh.

DRAM devices have a built-in counter for the refresh row address. By activating Column Address Strobe (CAS) before activating Row Address Strobe (RAS), this counter is selected to supply the row address instead of the address inputs.

CEC.

See *Core Event Controller*

circular buffer addressing.

The process by which the Data Address Generator (DAG) "wraps around" or repeatedly steps through a range of registers.

companding.

(compressing/expanding). The process of logarithmically encoding and decoding data to minimize the number of bits that must be sent.

conditional branches.

Jump or call/return instructions whose execution is based on testing an `if` condition.

core.

The processor, L1 memory, Event Controller, core timer, and Performance Monitoring registers.

Core Event Controller (CEC).

Part of the ADSP-BF561 Blackfin processor's two-stage event-control mechanism. The CEC works with the System Interrupt Controller (SIC) to prioritize and control all system interrupts. The CEC handles general-purpose interrupts and interrupts routed from the SIC.

CPLB.

See *Cacheability Protection Lookaside Buffer*

DAB.

See *DMA access bus*

DAG.

See *Data Address Generator*

Data Address Generator (DAG).

Processing component that provides memory addresses when data is transferred between memory and registers.

Data Register File.

A set of data registers that transfers data between the data buses and the computation units and provide local storage for operands.

data registers (Dreg).

Registers located in the computational units that hold operands for multiplier, ALU, or shifter operations.

delayed branch.

Jump or call/return instruction with the delayed branches modifier. In delayed branches, two (instead of four) instruction cycles are lost in the pipeline, because the processor executes the two instructions after the branch while the pipeline fills with instructions from the new branch.

descriptor block, DMA.

A set of parameters used by the direct memory access (DMA) controller to describe a set of DMA sequences. When successive DMA sequences are needed, the descriptor blocks are chained together so the completion of one DMA sequence autoinitiates and starts the next sequence.

descriptor loading, DMA.

The process in which the direct memory access (DMA) controller downloads a DMA descriptor from data memory and autoinitializes the DMA parameter registers.

DFT (Design For Testability).

A set of techniques that help designers of digital systems ensure that those systems will be testable.

Digital Signal Processor (DSP).

An integrated circuit designated for high-speed manipulation of analog information that has been converted into digital form.

direct branches.

Jump or call/return instructions that use absolute addresses that do not change at runtime (such as a program label), or they use a PC-relative address.

direct-mapped.

Cache architecture where each line has only one place that it can appear in the cache. Also described as 1-Way associative.

Direct Memory Access (DMA).

A way of moving data between system devices and memory in which the data is transferred through a DMA port without involving the processor (see *PIO*).

dirty, modified.

A state bit, stored along with the tag, indicating whether the data in the data cache line has been changed since it was copied from the source memory and, therefore, needs to be updated in that source memory.

DMA.

See *Direct Memory Access*

DMA Access Bus (DAB).

A peripheral bus that allows DMA-capable peripherals to gain access to memory.

DMA chaining.

The linking or chaining of multiple direct memory access (DMA) sequences. In chained DMA, the I/O processor loads the next DMA descriptor into the DMA parameter registers when the current DMA finishes and autoinitializes the next DMA sequence.

DMA descriptor registers.

Registers that hold the setup information for a direct memory access (DMA) process.

DPMC (Dynamic Power Management Controller).

A processor's control block that allows the user to dynamically control the processor's performance characteristics and power dissipation.

DQM Data I/O Mask Function.

The SDQM[3:0] pins provide a byte-masking capability on 8- or 16-bit writes to SDRAM. The SDQM[3:0] pins are not used to mask data on read cycles. For 16-bit wide SDRAM, only SDQM[1:0] are needed for byte masking.

DRAM (Dynamic Random Access Memory).

A type of semiconductor memory in which the data is stored as electrical charges in an array of cells, each consisting of a capacitor and a transistor. The cells are arranged on a chip in a grid of rows and columns. Since the capacitors discharge gradually—and the cells lose their information—the array of cells has to be refreshed periodically.

DSP.

See *Digital Signal Processor*

EAB.

See *External Access Bus*

EBC.

See *External Bus Controller*

EBIU.

See *External Bus Interface Unit*

EMB.

See *External Mastered Bus*

EPROM (Erasable Programmable Read-Only Memory).

A type of semiconductor memory in which the data is stored as electrical charges in isolated ("floating") transistor gates that retain their charges almost indefinitely without an external power supply. An EPROM is programmed by "injecting" charge into the floating gates—a process that requires relatively high voltage (usually 12V – 25V). Ultraviolet light, applied to the chip's surface through a quartz window in the package, will discharge the floating gates, allowing the chip to be reprogrammed.

EVT (Event Vector Table).

A table stored in low memory that contains sixteen 32-bit entries; each entry contains a vector address for an interrupt service routine (ISR). When an event occurs, instruction fetch starts at the address location in the corresponding EVT entry. See *ISR*.

exclusive, clean.

The state of a data cache line indicating the line is valid and the data contained in the line matches that in the source memory. The data in a clean cache line does not need to be written to source memory before it is replaced.

External Bus Interface Unit (EBIU).

A component that provides glueless interfaces to external memories. It services requests for external memory from the core, from the PCI bridge, or from a DMA channel.

External Access Bus (EAB).

A line or bus that allows the processor core and Memory DMA controller to directly access off-chip memory and PCI memory space to perform

instruction fetches, data loads, data stores, and high throughput memory-to-memory DMA transfers.

External Bus Controller (EBC).

A component that provides arbitration between the External Access Bus (EAB) and the External Mastered Bus (EMB), granting at most one requester per cycle.

External Mastered Bus (EMB).

A line or bus that allows the PCI to directly access internal L2 memory, off-chip memory, and the system memory-mapped registers (MMRs).

edge-sensitive interrupt.

A signal or interrupt the processor detects if the input signal is high (inactive) on one cycle and low (active) on the next cycle when sampled on the rising edge of `CLKIN`.

Endian format.

The ordering of bytes in a multibyte number. The processor uses big endian format—moves data starting with most significant bit (MSB) and finishing with least significant bit (LSB)—in almost all instances. The two exceptions are bit order for data transfer through the serial port and word order for packing through the external port. For compatibility with little endian (least significant first) peripherals, the processor supports both big and little endian bit order data transfers. Also for compatibility with little endian hosts, the processor supports both big and little endian word order data transfers.

external port.

A channel or port that extends the processor's internal address and data buses off-chip, providing the processor's interface to off-chip memory and peripherals.

FFT (Fast Fourier Transform).

An algorithm for computing the Fourier transform of a set of discrete data values. The FFT expresses a finite set of data points, for example a periodic sampling of a real-world signal, in terms of its component frequencies. Or conversely, the FFT reconstructs a signal from the frequency data. The FFT can also be used to multiply two polynomials. (For Jean Baptiste Joseph Fourier, 1768 – 1830, a French mathematician.)

FIFO (First In, First Out).

A hardware buffer or data structure from which items are taken out in the same order they were put in. Also known as a "shelf" from the analogy of pushing items onto one end of a shelf so that they fall off the other.

flag update.

A refresh from the processor to the status flags that occurs at the end of the cycle in which the status is generated. The flag update is available on the next cycle.

flash memory.

A type of single transistor cell, erasable memory in which erasing can only be done in blocks or for the entire chip.

Field Programmable Gate Array (FPGA).

A digital logic chip that can be programmed after production and that contains many thousands of gates.

FPGA.

See *Field Programmable Gate Array*.

fully associative.

Cache architecture where each line can be placed anywhere in the cache.

glueless.

No external hardware is required.

GSM (Global System for Mobile Communications).

A digital mobile telephone system that is widely used in Europe and other parts of the world. GSM uses a variation of Time Division Multiple Access (TDMA) and operates at either the 900 MHz or 1800 MHz frequency band.

Harvard architecture.

A processor memory architecture that uses separate buses for program and data storage. The two buses let the processor get a data word and an instruction simultaneously.

HLL (High Level Language).

A programming language that provides some level of abstraction above assembly language, often using English-like statements, where each command or statement corresponds to several machine instructions.

IDLE.

An instruction that causes the processor to cease operations, holding its current state until an interrupt occurs. Then, the processor services the interrupt and continues normal execution.

IEEE (Institute of Electrical and Electronics Engineers).

An international technical professional society based in the United States.

index.

Address portion that is used to select an array element (for example, line index).

Index registers.

A Data Address Generator (DAG) register that holds an address and acts as a pointer to memory.

indirect branches.

Jump or call/return instructions that use a dynamic address from the data address generator, evaluated at runtime.

I/O processor.

Direct memory access (DMA) controllers, DMA channel arbitration, and peripheral-to-bus connections that handle input and output operations relieving the burden on the processor. The processor uses a distributed DMA architecture with DMA controllers for each DMA-capable peripheral. Most ports have some direct (non-DMA) access to internal memory and I/O memory.

I/O processor register.

One of the control, status, or data buffer registers of the processor's on-chip I/O processor.

input clock.

Device that generates a steady stream of timing signals to provide the frequency, duty cycle, and stability to allow accurate internal clock multiplication via the phase locked loop (PLL) module.

internal memory bank.

There are several internal memory banks on a given SDRAM row. An internal bank in a specific row cannot be activated (opened) until the previous internal bank in that row has been precharged.

The SDRAM controller (SDC) does not support interleaved accesses. The bank address can be thought of as part of the row address. The SDC also assumes that all SDRAMs to which it interfaces have four internal banks.

Do not confuse these “SDRAM internal banks,” which are internal to the SDRAM and are selected with the bank address, with the four “SDRAM banks” or “external banks” that are enabled by the SMS[3:0] pins.

interrupt.

An event that suspends normal processing and temporarily diverts the flow of control through an interrupt service routine (ISR). See *ISR*.

invalid.

Describes the state of a cache line. When a cache line is invalid, a cache line match cannot occur.

IrDA (Infrared Data Association).

A nonprofit trade association that established standards for ensuring the quality and interoperability of devices using the infrared spectrum.

isochronous.

Processes where data must be delivered within certain time constraints.

ISR (Interrupt Service Routine).

Software that is executed when a specific interrupt occurs. A table stored in low memory contains pointers, also called vectors, that direct the processor to the corresponding ISR. See *EVT*.

JEDEC.

(formerly known as the Joint Electronic Device Engineering Council, now called the JEDEC Solid State Technology Association) The semiconductor engineering standardization body of the Electronic Industries Alliance (EIA), an electronic industry trade association.

JTAG (Joint Test Action Group).

An IEEE Standards working group that defines the IEEE 1149.1 standard for a test access port for testing electronic devices.

JTAG port.

A channel or port that supports the IEEE standard 1149.1 JTAG standard for system test. This standard defines a method for serially scanning the I/O status of each component in a system.

jump.

A permanent transfer of the program flow to another part of program memory.

latency.

The overhead time used to find the correct place for memory access and preparing to access it.

Least Recently Used algorithm.

Replacement algorithm used by cache that first replaces lines that have been unused for the longest time.

Least Significant Bit (LSB).

The last or rightmost bit in the normal representation of a binary number—the bit of a binary number giving the number of ones.

Length registers.

A Data Address Generator (DAG) register that sets up the range of addresses in a circular buffer.

Level 1 (L1) memory.

Memory that is directly accessed by the core with no intervening memory subsystems between it and the core.

Level 2 (L2) memory.

Memory that is at least one level removed from the core. L2 memory has a larger capacity than L1 memory, but it requires additional latency to access.

level-sensitive interrupts.

A signal or interrupt that the processor detects if the input signal is low (active) when sampled on the rising edge of `CLKIN`.

LIFO (Last In, First Out).

A data structure from which the next item taken out is the most recent item put in. Also known as a "stack" from the analogy of stacked plates in a cafeteria in which the most recent plate placed on the stack is on top and thus will be the next plate removed. See *Stack*.

little endian.

The native data store format of the processor. Words and half words are stored in memory (and registers) with the least significant byte at the lowest byte address and the most significant byte at the highest byte address of the data storage location.

loop.

A sequence of instructions that executes several times with zero overhead.

LRU (Least Recently Used algorithm).

A rule based on the observation that, in general, the cache entry that has not been accessed for longest is least likely to be accessed in the near future.

LSB.

See *Least Significant Bit*.

MAC (Multiply/Accumulate).

A floating-point operation that multiplies two numbers and then adds a third to get the result (see *Multiply Accumulator*).

memory bank.

A unit of the processor's external memory space, which may be addressed by either data address generator. External memory banks also may be configured for size and access waitstates.

memory block.

A unit of the processor's internal memory space, which is associated with a Data Address Generator (DAG).

Memory Management Unit (MMU).

A component of the processor that supports protection and selective caching of memory by using Cacheability Protection Lookaside Buffers (CPLBs).

Mode Register.

SDRAM devices contain internal configuration registers which allows specification of the SDRAM device's functionality. After powerup and before executing a read or write to the SDRAM memory space, the application must trigger the SDRAM controller (SDC) to write the SDRAM's Mode register. The write of the SDRAM's Mode register is triggered by writing a 1 to the PSSE bit in the SDRAM Memory Global Control register (EBIU_SDGCTL), and then issuing a read or write transfer to the SDRAM address space. The initial read or write triggers the SDRAM powerup sequence to be run, which programs the SDRAM's Mode register with the CAS latency from the EBIU_SDGCTL register. This initial read or write to SDRAM takes many cycles to complete. Note for most applications the SDRAM powerup sequence and writing of the Mode register needs to be done only once. Once the powerup sequence has completed, the PSSE bit should not be set again unless a change to the Mode register is desired.

modified addressing.

The process whereby the Data Address Generator (DAG) produces an address that is incremented by a value or a register.

Modify address.

The process by which a Data Address Generator (DAG) increments the stored address without performing a data move.

Modify register.

A Data Address Generator (DAG) register that provides the increment or step size by which an index register is pre- or post-modified during a register move.

MMR (Memory-Mapped Register).

A location in main memory used by the processor as if it were a register. Registers are high-speed memory locations in the processor that can be addressed with just a few bits; memory-mapped registers must be addressed in the same way as other memory locations and the speed of MMR accesses is limited to the speed of the system's data bus.

MMU.

See *Memory Management Unit*

MSB (Most Significant Bit).

The first or leftmost bit in the normal representation of a binary number—the bit of a binary number with the greatest weight (2(n-1)).

multifunction computations.

The parallel execution of multiple computational instructions. These instructions complete in a single cycle, and they combine parallel operation of the computational units and memory accesses. The multiple

operations perform the same as if they were in corresponding single function computations.

multiplier.

A computational unit that performs fixed-point multiplication and executes fixed-point multiply/add and multiply/subtract operations.

Multiply Accumulator.

An accumulator register in which operands are multiplied and added to the contents.

NMI (Nonmaskable Interrupt).

A high priority interrupt that cannot be disabled by another interrupt.

NRZ (Non-Return-to-Zero).

A binary encoding scheme in which a 1 is represented by a change in the signal and a 0 by no change—there is no return to a reference (0) voltage between encoded bits. This method eliminates the need for a clock signal.

NRZI (Non-Return-to-Zero Inverted).

A binary encoding scheme in which a 0 is represented by a change in the signal and a 1 is represented by no change—there is no return to a reference (0) voltage between encoded bits. This method eliminates the need for a clock signal.

orthogonal.

The characteristic of being independent. An orthogonal instruction set allows any register to be used in an instruction that references a register.

PAB.

See *Peripheral Access Bus*

page size.

The amount of memory which has the same row address and can be accessed with successive read or write commands without needing to activate another row. The page size can be calculated for 32-bit and 16-bit SDRAM banks with these formulas:

- 32-bit SDRAM banks: page size = 2(CAW + 2)
- where `CAW` is the column address width of the SDRAM, plus 2 because the SDRAM bank is 32 bits wide (2 address bits = 4 bytes).
- 16-bit SDRAM banks: page size = 2(CAW + 1)
- where `CAW` is the column address width of the SDRAM, plus 1 because the SDRAM bank is 16 bits wide (1 address bit = 2 bytes).

PC (Program Counter).

A register that contains the address of the next instruction to be executed. The PC is automatically incremented after each instruction is fetched to point to the following instruction.

PCB (Printed Circuit Board).

The circuit board on which the processor device is mounted. The PCB provides power and signal interfaces to the processor, as well as mechanical support and cooling.

PCI.

See *Peripheral Component Interconnect*

peripheral.

The processor's peripherals include timers, Real-Time Clock, USB, programmable flags, UARTs, SPORTs, and SPIs.

Peripheral Access Bus (PAB).

A low and predictable latency bus that keeps core stalls to a minimum and allows for manageable interrupt latencies to time critical peripherals.

Peripheral Component Interconnect (PCI).

A bus system that provides a bus bridge between the External Access Bus (EAB) and the External Mastered Bus (EMB) internal to the processor and an external PCI bus.

PF (Programmable Flag).

General-purpose I/O pins. Each PF pin can be individually configured as either an input or an output pin, and each PF pin can be further configured to generate an interrupt.

Phase Locked Loop (PLL).

An on-chip frequency synthesizer that produces a full speed master clock from a lower frequency input clock signal.

PIO (Programmed Input/Output).

A way of moving data between system devices and memory in which all data must pass through the processor (see *Direct Memory Access*).

PLL.

See *Phase Locked Loop*.

precision.

The number of bits after the binary point in the storage format for the number.

post-modify addressing.

The process in which the Data Address Generator (DAG) provides an address during a data move and autoincrements the stored address for the next move.

Precharge command.

The Precharge command closes a specific internal bank in the active page or all internal banks in the page. The SDRAM Controller (SDC) always does a Precharge All, closing all internal banks.

pre-modify addressing.

The process in which the Data Address Generator (DAG) provides an address during a data move without incrementing the stored address.

PWM (Pulse Width Modulation).

Also called Pulse Duration Modulation (PDM), PWM is a pulse modulation technique in which the duration of the pulses is varied by the modulating voltage.

RAS (Row Address Strobe).

A signal sent from the Memory Management Unit (MMU) to a DRAM device to indicate validity of row address lines.

Real-Time Clock (RTC).

A component that generates timing pulses for the digital watch features of the processor, including time of day, alarm, and stopwatch countdown features.

replacement policy.

The function used by the processor to determine which line to replace on a cache miss. Often, a least recently used (LRU) algorithm is employed.

ROM (Read-Only Memory).

A data storage device manufactured with fixed contents. This term is most often used to refer to non-volatile semiconductor memory.

RTC.

See *Real-Time Clock*

RZ (Return-to-Zero modulation).

A binary encoding scheme in which two signal pulses are used for every bit. A 0 is represented by a change from the low voltage level to the high voltage level; a 1 is represented by a change from the high voltage level to the low voltage level. A return to a reference (0) voltage is made between encoded bits.

RZI (Return-to-Zero-Inverted modulation).

A binary encoding scheme in which two signal pulses are used for every bit. A 1 is represented by a change from the low voltage level to the high voltage level; a 0 is represented by a change from the high voltage level to the low voltage level. A return to a reference (0) voltage is made between encoded bits.

saturation (ALU saturation mode).

A state in which all positive fixed-point overflows return the maximum positive fixed-point number, and all negative overflows return the maximum negative number.

SDRAM (Synchronous Dynamic Random Access Memory).

A form of DRAM that includes a clock signal with its other control signals. This clock signal allows SDRAM devices to support “burst” access modes that clock out a series of successive bits.

SDRAM banks.

Four regions of memory that can be configured to be 16 MB, 32 MB, 64 MB, or 128 MB and are selected by the SMS[3:0] pins. Each bank can be selected to be either all 32 bits wide or all 16 bits wide.

Do not confuse the “SDRAM internal banks” which are internal to the SDRAM and are selected with the bank address, with the “SDRAM banks” or “external banks” that are enabled by the SMS[3:0] pins.

SDRAM DIMMs.

Dual In-line Memory Modules, or DIMMs, are an industry-standard SDRAM packaging option. The DIMM consists of a small, standardized form factor board, populated with SDRAM devices on one or both sides. DIMMs are populated with sufficient SDRAM to provide either 32-bit or 64-bit data paths, with some configurations supporting additional data bits for parity protection.

Many DIMMs also include a serial presence detect port, which provides access to an on-DIMM serial ROM that includes information describing the type and characteristics of SDRAMs on the DIMM. This information can be read to determine the type, size, and timing parameters of installed DIMMs at boot time. One of the Serial Peripheral Interface (SPI) or SPORT peripherals of the processor can be used to interface to the DIMM serial presence detect port, if dynamic boot time determination of SDRAM configuration is required.

Self-Refresh.

When the SDRAM is in Self-Refresh mode, the SDRAM’s internal timer initiates Auto-Refresh cycles periodically, without external control input. The SDRAM Controller (SDC) must issue a series of commands including the Self-Refresh command to put SDRAM into low power mode, and it must issue another series of commands to exit Self-Refresh mode. Entering Self-Refresh mode is programmed in the SDRAM Memory Global Control register (EBIU_SDGCTL) and any access to the SDRAM address

space causes the SDC to exit SDRAM from Self-Refresh mode. See "Entering and Exiting Self-Refresh Mode (SRFS)" on page 16-40.

Serial Peripheral Interface (SPI).

A synchronous serial protocol used to connect integrated circuits.

serial ports.

An input/output device on the processor. The processor uses two synchronous serial ports that provide inexpensive interfaces to a wide variety of digital and mixed-signal peripheral devices.

set.

A group of *N*-line storage locations in the Ways of an *N*-Way cache, selected by the Index field of the address.

set associative.

Cache architecture that limits line placement to a number of sets (or Ways).

shifter.

A computational unit that completes logical and arithmetic shifts on 16-bit operands and derives exponents.

SIC (System Interrupt Controller).

Part of the processor's two-level event control mechanism. The SIC works with the Core Event Controller (CEC) to prioritize and control all system interrupts. The SIC provides mapping between the peripheral interrupt sources and the prioritized general-purpose interrupt inputs of the core.

SIMD (Single-Instruction, Multiple-Data).

A parallel computer architecture in which a collection of data is processed simultaneously under one instruction.

SP (Stack Pointer).

A register that points to the top of a stack. If the stack contains data, the SP points to the most recently pushed item; if the stack does not contain data, the SP points to the first empty location, where the next item will be pushed.

SPI.

See *Serial Peripheral Interface*

SRAM.

See *Static Random Access Memory*

stack.

A data structure for storing items that are to be accessed in last in, first out (LIFO) order. When a data item is added to the stack, it is "pushed"; when a data item is removed from the stack, it is "popped."

Static Random Access Memory (SRAM).

Very fast memory that does not require periodic refreshing.

synchronous process.

A process that runs only as a result of some other process being completed or as a result of a handing-off operation.

system.

The system includes the peripheral set (Timers, Real-Time Clock, programmable flags, UART, SPORTs, PPI, and SPIs), the external memory controller (EBIU), the Memory DMA controller, as well as the interfaces between these peripherals, and the optional, external (off-chip) resources.

System clock (SCLK).

A component that delivers pulses at a frequency determined by a programmable divider ratio from the core clock.

System Interrupt Controller (SIC).

Component that maps and routes events from peripheral interrupt sources to the prioritized, general-purpose interrupt inputs of the CIC.

subroutine.

A self-contained section of a program, with an entry point and an exit, that usually performs a single task. The main program flow branches to a subroutine when its task is required, and the subroutine returns (branches back) to the main flow when its task has been completed.

t_{RAS}.

Required delay between issuing a Bank Activate command and issuing a Precharge command, and between the Self-Refresh command and the exit from Self-Refresh. The `TRAS` bit field in the SDRAM Memory Global Control register (`EBIU_SDGCTL`) is 4 bits wide and can be programmed to be 1 to 15 clock cycles long. See "Selecting the Bank Activate Command Delay (TRAS)" on page 16-43.

t_{RP}.

Required delay between issuing a Precharge command and:

- issuing a Bank Activate command
- issuing an Auto-Refresh command
- issuing a Self-Refresh command

The `TRP` bit field in the SDRAM Memory Global Control register (`EBIU_SDGCTL`) is 3 bits wide and can be programmed to be 1 to 7 clock cycles long. See "Selecting the Precharge Delay (TRP)" on page 16-44.

t_{RCD}.

Required delay between a Bank Activate command and the start of the first Read or Write command. The TRCD bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is 3 bits wide and can be programmed to be from 1 to 7 clock cycles long.

t_{WR}.

Required delay between a Write command (driving write data) and a Precharge command. The TWR bit field in the SDRAM Memory Global Control register (EBIU_SDGCTL) is 2 bits wide and can be programmed to be from 1 to 3 clock cycles long.

t_{RC}.

Required delay between issuing successive Bank Activate commands to the same SDRAM internal bank. This delay is not directly programmable. The t_{RC} delay must be satisfied by programming the TRAS and TRP fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

t_{RFC}.

Required delay between issuing an Auto-Refresh command and a Bank Activate command and between issuing successive Auto-Refresh commands. This delay is not directly programmable and is assumed to be equal to t_{RC}. The t_{RC} delay must be satisfied by programming the TRAS and TRP fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

t_{XSR}

Required delay between exiting Self-Refresh mode and issuing the Auto-Refresh command. This delay is not directly programmable and is assumed to be equal to t_{RC}. The t_{RC} delay must be satisfied by programming the t_{RAS} and t_{RP} fields to ensure that $t_{RAS} + t_{RP} \geq t_{RC}$.

tag.

Upper address bits, stored along with the cached data line, to identify the specific address source in memory that the cached line represents.

TAP (Test Access Port).

See *JTAG port*

TDM.

See *Time Division Multiplexing*

tristate versus Three-state.

Analog Devices documentation uses the word *three-state* instead of *tristate*. Tristate™ is a trademarked term owned by National Semiconductor.

Time Division Multiplexing (TDM).

A method used for transmitting separate signals over a single channel. Transmission time is broken into segments, each of which carries one element. Each word belongs to the next consecutive channel so that, for example, a 24-word block of data contains one word for each of the 24 channels.

UART.

See *Universal Asynchronous Receiver Transmitter*

UDC (USB Device Controller).

Part of the Universal Serial Bus Device (USBD), a module that implements the low level USB protocol. The UDC manages interaction with the USB host using the USB serial link, and presents data and command transactions to the application.

Universal Serial Bus (USB).

A module that handles USB protocol requirements. It also contains hardware to connect it to the processor's peripheral buses and to support an operating system.

Universal Asynchronous Receiver Transmitter (UART).

A module that contains both the receiving and transmitting circuits required for asynchronous serial communication.

USB.

See *Universal Serial Bus*

USBD (Universal Serial Bus Device).

The module in the processor that controls the Universal Serial Bus (USB). The USBD consists of the USB Device Controller (UDC) core module and a front-end interface. The USBD connects the DMA Access Bus (DAB) and the Peripheral Access Bus (PAB) with and off-chip USB transceiver.

Valid.

A state bit (stored along with the tag) that indicates the corresponding tag and data are current and correct and can be used to satisfy memory access requests.

via.

A circuit connection between layers of a PCB.

victim.

A dirty cache line that must be written to memory before it can be replaced to free space for a cache line allocation.

Von Neumann architecture.

The architecture used by most non-DSP microprocessors. This architecture uses a single address and data bus for memory access. (For John von Neumann, 1903-1957, a Hungarian-born mathematician.)

Way.

An array of line storage elements in an *N*-Way cache.

W1C.

See *Write-1-to-Clear*

W1S.

See *Write-1-to-Set*

Write-1-to-Clear (W1C) bit.

A control or status bit that can be cleared (= 0) by being written to with 1.

Write-1-to-Set (W1S) bit.

A control or status bit that is set by writing 1 to it. It cannot be cleared by writing 0 to it.

write back.

A cache write policy (also known as copyback). The write data is written only to the cache line. The modified cache line is written to source memory only when it is replaced.

write through.

A cache write policy (also known as store through). The write data is written to both the cache line and to source memory. The modified cache line is *not* written to the source memory when it is replaced.

I INDEX

Symbols

μ-law companding, 12-2, 12-33, 12-57

Numerics

0x prefix, xliii
16-bit booting, 17-2
16-bit operations, 2-34
16-bit SDRAM bank, 16-56
256-lead MBGA package, 17-1
32-bit operations, 2-27, 2-28
8-bit booting, 17-2

A

A10 pin of SDRAM, 16-34
AC (address calculation), 4-7
access
  request priority. *See also* arbitration
  DMA stalls to L1 or L2 memory, 7-13
  IMDMA, 7-11
  latency and throughput for L2 memory, 7-12
  priority, L2 port, 7-11
accumulator result registers A[1: 0], 2-6, 2-37, 2-43
active low/high frame syncs, serial port, 12-37
active mode, 1-22, 8-14
active video only mode, PPI, 11-19
address calculation (AC), 4-7
addressing
  See also auto-decrement; auto-increment; bit-reversed; circular-buffer; indexed; indirect; modified; post-increment; post-modify; pre-modify; data address generators
  modes, 5-15
address-tag compare operation, 6-14
A-law companding, 12-2, 12-33, 12-57
alignment exceptions, 6-72
allocating system stack, 4-62
alternate frame sync mode, 12-40
alternate timing, serial port, 12-39
ALU. *See* arithmetic logic unit (ALU)
AMC
  EBIU block diagram, 16-3
  SCLK[1] function, 16-39
  timing parameters, 16-12
AMS, 16-9
Analog Devices product information, xli
AND, logical, 2-23
arbitration, 7-21
  DMA bus, 7-9, 7-17
  EBIU, EAB bus and DMA bus access, 7-21
  EBIU access, 7-21
  PAB, 7-14
  priority for Sys L2 port access request, 7-12
ARDY, 16-13, 16-19
arithmetic formats summary, 2-15, 2-16

arithmetic logic unit (ALU), 1-3, 2-1, 2-23, 2-36
  arithmetic, 2-13
  data flow, 2-33
  data types, 2-12
  functions, 2-23
  inputs and outputs, 2-25
  instructions, 2-25, 2-28, 2-36
  operations, 2-23, 2-28
  status, 2-22
arithmetic operations, 2-23
arithmetic shift (ASHIFT) instruction, 2-14, 2-48
arithmetic shifts, 2-1
arithmetic status register (ASTAT), 2-23
ASIC/FPGA designs, 16-1
assembly language, 2-1
ASTAT (arithmetic status register), 2-23
asynchronous controller, 1-13
asynchronous interfaces supported, 16-1
asynchronous memory
  interface, 16-8
  region, 16-2
asynchronous memory bank address range (table), 16-9
asynchronous memory bank control registers (EBIU_AMBCTLx), 16-12
asynchronous memory controller. *See* AMC
asynchronous memory global control register (EBIU_AMGCTL), 16-9
asynchronous memory interfaces, example of, 17-9
asynchronous memory interface signals (table), 16-5
asynchronous read, 16-17, 16-19
asynchronous read bus cycle (figure), 16-18
asynchronous serial communications, 13-2
asynchronous write, 16-19
asynchronous write and read bus cycle (figure), 16-19
ASYNC memory banks, 16-4
atomic operations, 6-73, 17-7
autobaud detection, 13-1, 15-38
auto-decrement addressing, 5-10
auto-increment addressing, 5-10, 5-14
auto-refresh
  command, 16-27, 16-65
  description, 16-59
  timing, 16-50
avoiding bus contention, 16-13

B

bandwidth, EAB bus, 7-22
bank activate command, 16-26, 16-62
  selecting delay, 16-43
bank address, 16-48, 16-49
bank size, 16-48, 16-49
bank size encodings (table), 16-53, 16-54
bank width, 16-48, 16-49
barrel-shifter. *See* shifter
base registers (B[3:0]), 2-7, 5-2, 5-7
baud rate, SPI booting, 3-21
B (base) registers, 2-7, 5-7
biased rounding, 2-18
binary multiplication, C-4
binary numbers, 2-3
binary semaphores, 17-7
bit manipulation, 2-1
  bit clear, 2-52
  bit set, 2-52
  bit test, 2-52
  bit toggle, 2-52
bit order, selecting, 12-32
bit-reversed addressing, 5-1, 5-9
bits marked x, xlvi
Blackfin family, core architecture, 1-1
block diagram, 1-2
block diagram, core, 7-5

block diagrams
interrupt processing, 4-22
PLL, 8-3
SDRAM, 16-32
SPI, 10-2
SPORT, 12-5
block floating point format, C-6
BMODE
bits, 3-14
pins, 4-42
state, 3-13
BMODE (boot mode) pins, 3-21
booting from a ROM, 17-2
booting from external 16-bit memory, 17-2
booting the processor, 17-2
boot kernel, 3-19
boot mode, 1-24
boot ROM, loading user code, 3-19
b# prefix, xliii
branch, 4-9
branch, conditional, 4-13
branch latency, 4-10
conditional branches, 4-14
unconditional branches, 4-14
branch prediction, 4-14
branch target, 4-12
branch target address, unconditional branches, 4-14
B-registers (base), 2-7, 5-2, 5-7
broadcast mode, 10-2, 10-14, 10-22
buffers, cacheability protection lookaside buffers (CPLBs), 6-47, 6-48
burst length, 16-26, 16-58
burst stop command, 16-26
burst type, 16-26
bus access arbitration, 7-21
bus arbitration. *See* arbitration
bus connection and arbitration, DMA, 7-17
bus contention, avoiding, 16-13
bus contention on the external memory bus, how to avoid, 17-9
bus error, EBIU, 16-8
bus priorities, 17-8
bypass capacitor placement, 17-14
BYPASS field, 8-7
byte address, 16-48, 16-49
byte order, 2-12

C

cache, 6-12
coherency support, 6-72
mapping into data banks, 6-31
validity of cache lines, 6-14
cacheability protection lookaside buffers (CPLBs), 3-5, 3-11, 6-12, 6-47, 6-48
cache block (definition), 6-75
cache hit, 6-75
address-tag compare, 6-14
data cache access, 6-34
definition, 6-14
processing, 6-15
cache inhibited accesses, 6-73
cache line
components, 6-12
definition, 6-75
states, 6-34
cache miss, 6-75
definition, 6-34
replacement policy, 6-16
CALL instruction, 4-9, 4-11
range, 4-11
capacitive loads, 16-24
capacitors, decoupling, 17-13
CAS before RAS, 16-27
CAS latency, 16-27
selecting, 16-42
CAW, 16-48, 16-49
CBR refresh, 16-27
CC flag bit, 4-12

CCIR-656. *See* ITU-656
CCITT G.711 specification, 12-33
CCLK (core clock), 8-4
  derivation, 8-1
  status by operating mode, 8-12
CC status bit, 4-10
channels
  defined, serial, 12-55
  serial port TDM, 12-56
  serial select offset, 12-55
CHNL bit, 12-54
circuit board design and layout, 17-12
circular buffer addressing, 5-6
  registers, 5-6
  wraparound, 5-8
CKELOW bit, 8-25
CL bit, 16-27, 16-36
clean (definition), 6-75
CLKIN (input clock), 8-1, 8-2
CLKIN to CCLK, changing the multiplier, 8-19
clock
  core processor, 7-5
  EBIU, 16-1
  frequency for SPORT, 12-27
  internal, 7-5
  load, 16-39
  managing clocks, 17-5
  setting up, 12-31
  types, 17-5
clock divide modulus register, 12-27
clock input (CLKIN) pin, 17-5
clock phase, SPI, 10-20, 10-22
clock polarity, SPI, 10-20
clock rate, SPORT, 12-2
code examples
  active mode to full-on mode, 8-21
  CSYNC, 6-75
  epilog code for nested ISR, 4-55
  execution trace recreation, 18-18
code examples *(continued)*
  full-on mode to active mode, 8-22
  interrupt enabling and disabling, 6-74
  load base of MMRs, 6-74
  loop, 4-16
  modification of PLL, 8-19
  prolog code for nested ISR, 4-55
  restoration of the control register, 6-74
code patching, 18-4
column address, 16-48, 16-49
column address strobe latency, 16-27
command inhibit command, 16-67
commands
  auto-refresh, 16-25, 16-27, 16-30, 16-31, 16-34, 16-40, 16-44, 16-50, 16-59, 16-65
  bank activate, 16-26, 16-30, 16-31, 16-34, 16-43, 16-44, 16-51, 16-58, 16-62
  burst stop, 16-26
  command inhibit, 16-67
  load mode register, 16-36, 16-64
  no operation, 16-67
  parallel refresh, 16-33
  precharge, 16-29, 16-30, 16-31, 16-34, 16-37, 16-40, 16-43, 16-44, 16-45, 16-58, 16-62
  precharge all, 16-36
  read, 16-30, 16-34, 16-42, 16-58
  read/write, 16-65
  SDC, 16-61
  self-refresh, 16-25, 16-29, 16-30, 16-40, 16-65
  transfer initiate, 10-24
  write, 16-30, 16-31, 16-34, 16-45, 16-58
companding, 12-47, 12-57
  defined, 12-33
  lengths supported, 12-33
  multichannel operations, 12-57

computation
instructions, 2-1
status, 2-22
computational units, 2-1, 2-55
conditional
branches, 6-70, 6-71
instructions, 2-22
JUMP instruction, 4-10
conditional branches, 4-13, 4-14
branch latency, 4-14
conditional instructions, 4-3
condition code (CC) flag bit, 4-12
configuration
L1 instruction memory, 6-11
L1 SRAM, 6-2
precautions before changing, 6-10
SDRAM, 16-24
SPORT, 12-10
connection and arbitration, DMA bus, 7-17
content-addressable memory (CAM), 6-47
contention, bus, avoiding, 16-13
controller, DMA, 7-2, 7-18
control/optimization, DMA traffic, 7-13
control register
data memory, 6-24
EBIU, 16-7
instruction memory, 6-7
PLL, 8-3
restoration, 6-74
conventions, xliii
convergent rounding, 2-18
coordinate between multiple processors, 17-6
core
access to flag configuration, 14-8
clock/system clock ratio control, 8-4
double-fault condition, 4-41
core architecture, 1-3, 1-5
core block diagram, 7-5
core clock (CCLK), 8-4, 17-5
core clock divider ratio, 17-6
core double-fault reset, 3-13
core event
in EVT, 4-40
MMR location, 4-40
core event controller (CEC), 1-10, 4-18
core event vector table (table), 4-40
core interfaces, DSP, 7-5
core interrupt latch register (ILAT), 4-37
core interrupt mask register (IMASK), 4-36
core interrupts, 17-6
core interrupts pending register (IPEND), 3-2, 4-38
core-only software reset, 3-13, 3-18, 3-20
core timer control register (TCNTL), 15-50
core timer count register (TCOUNT), 15-51
core timer period register (TPERIOD), 15-52
core timer scale register (TSCALE), 15-53
counter cycle, 4-4, 18-23
cross options, 2-34
crosstalk, 17-13
crystal oscillator, 17-5
CSYNC, 6-70
code example, 6-75
current address registers
(DMAx_CURR_ADDR), 9-20
(MDMA_yy_CURR_ADDR), 9-20
current descriptor pointer registers
(DMAx_CURR_DESC_PTR), 9-19
(MDMA_yy_CURR_DESC_PTR), 9-19
current inner loop count registers
(DMAx_CURR_X_COUNT), 9-21
(MDMA_yy_CURR_X_COUNT), 9-21

current outer loop count registers
(DMAx_CURR_Y_COUNT), 9-22
(MDMA_yy_CURR_Y_COUNT), 9-22
current spikes, reducing noise-producing, 17-12
customer support, xl
cycle counters, 4-4, 18-23, 18-24
CYCLES and CYCLES2 (execution cycle count registers), 18-24

D

DAB, DMA access bus, 7-4, 7-16, 7-17
DAG. *See* data address generators
data cache control instructions, 6-38
data delays, 17-8
data-driven interrupts, 9-27
data flow, 2-1
data formats, 2-3, 2-4, 2-11
SPORT, 12-32
data I/O mask pins, 16-33
data mask encodings, 16-57
data memory, L1, 6-24
data memory control register (DMEM_CONTROL), 6-24, 6-48
data move, serial port operations, 12-42
data operations, CPLB, 6-48
data overflow, 12-35
data register file, 2-5, 2-6
data registers, 2-5, 3-4
data sampling, serial, 12-37
data SRAM, L1, 6-26
data store format, 6-76
data test command register (DTEST_COMMAND), 6-41
data test data registers (DTEST_DATAx), 6-42
data transfer latency
DMA, 7-13
EAB, 7-22
data transfers
data register file, 2-6
SPI, 10-2
using PPI, 11-33
data types, 2-10, 2-21
data underflow, 12-35
data watchpoint address control register (WPDACTL), 18-12
data watchpoint address count value registers (WPDACNTn), 18-11
data watchpoint address registers (WPDAn), 18-11
data word
serial data formats, 12-20
date address generators (DAGs), 3-11, 5-1, 5-21
addressing modes, 5-15
DAG0 CPLB miss, 4-48
DAG0 misaligned access, 4-48
DAG0 multiple CPLB hits, 4-48
DAG0 protection violation, 4-48
DAG1 CPLB miss, 4-48
DAG1 misaligned access, 4-48
DAG1 multiple CPLB hits, 4-48
DAG1 protection violation, 4-48
instructions, 5-13, 5-16
register modification, 5-12
registers, 2-5, 2-7
support for branches, 4-3
DCB
DCB1, 7-8
DCB2, 7-9
DCB3, 7-9
DCB4, 7-9
DCB, DMA core bus, 7-4, 7-16, 7-17
DCB bus arbitration, 7-18
DCBS (L1 data cache bank select) bit, 6-32
DCPLB address registers (DCPLB_ADDRx), 6-60

DCPLB_ADDRx (DCPLB address registers), 6-60
DCPLB data registers (DCPLB_DATAx), 6-57
DCPLB_DATAx (DCPLB data registers), 6-57
DCPLB_FAULT_ADDR (DCPLB fault address register), 6-64
DCPLB fault address register (DCPLB_FAULT_ADDR), 6-64
DCPLB_STATUS (DCPLB status register), 6-63
DCPLB status register (DCPLB_STATUS), 6-63
DEB, DMA external bus, 7-4, 7-16, 7-17
debug features, 18-1
DEC (instruction decode), 4-7
decoupling capacitor recommendations, 17-13
deep-sleep mode, 1-23, 8-14
deferring exception processing, 4-57
delay count register (PPI_DELAY), 11-11
design considerations, high frequency, 17-11
DF (divide frequency), 8-3, 8-9
DI_EN bit, 9-14
direct branch, 4-10
direct mapped (definition), 6-75
direct memory access (DMA), 1-11, 9-1, 9-40
dirty (definition), 6-75
disable interrupts (CLI) instruction, 6-75
disabling
  interrupts, global, 4-39
  PLL, 8-17
DISALGNEXPT instruction, 5-13
DI_SEL bit, 9-14
DITFS bit, 12-15, 12-25, 12-41
divide frequency (DF) bit, 8-3, 8-9
divide primitives (DIVS, DIVQ), 2-12, 2-36
DIVQ instruction, 2-36
DIVS instruction, 2-36
DMA, 9-1, 9-40
  and SPORT, 12-3
  and UART, 13-16
  autobuffer mode, 9-28, 9-37
  buffer size, multichannel SPORT, 12-60
  channel registers, 9-28
  channels, 9-53
  continuous transfers using autobuffering, 9-56
  descriptor array, 9-35
  descriptor elements, 9-7
  descriptor lists, 9-36
  descriptor queue management, 9-59
  descriptor structures, 9-58
  direction, 9-15
  DMA error interrupt, 9-28
  flex descriptor structure, 9-28
  flow chart, 9-31
  for SPI transmit, 10-16, 10-32
  memory DMA, 9-40
  operation flow, 9-31
  polling registers, 9-54
  refresh, 9-35
  register naming conventions, 9-6
  serial port block transfers, 12-42
  single-buffer transfers, 9-55
  software management, 9-53
  startup, 9-33
  stopping, 9-37
  synchronization, 9-53
  two-dimensional (2D), 9-39
  with PPI, 11-31
DMA, related buses, 7-16
DMA2D bit, 9-15
DMA access bus (DAB), 7-4, 7-16, 7-17

DMA access to L1 or L2 memory, stalls, 7-13
DMA bus
  access arbitration, 7-21
  arbitration, 7-9
  connection and arbitration, 7-17
DMA configuration registers
  (DMAx_CONFIG), 9-12
  (MDMA_yy_CONFIG), 9-12
DMA controller, 1-2, 7-2, 7-18
DMA core bus (DCB), 7-4, 7-16, 7-17
DMA data transfer latency, 7-13
DMA_EN bit, 9-15
DMA external bus (DEB), 7-4, 7-16, 7-17
DMA traffic control/optimization, 7-13
DMAx_CONFIG (DMA configuration registers), 9-12
DMAx_CURR_ADDR (current address registers), 9-20
DMAx_CURR_DESC_PTR (current descriptor pointer registers), 9-19
DMAx_CURR_X_COUNT (current inner loop count registers), 9-21
DMAx_CURR_Y_COUNT (current outer loop count registers), 9-22
DMAx_IRQ_STATUS (interrupt status registers), 9-26
DMAx_NEXT_DESC_PTR (next descriptor pointer registers), 9-10
DMAx_PERIPHERAL_MAP (peripheral map registers), 9-23
DMAx_START_ADDR (start address registers), 9-11
DMAx_X_COUNT (inner loop count registers), 9-16
DMAx_X_MODIFY (inner loop address increment registers), 9-16
DMAx_Y_COUNT (outer loop count registers), 9-17
DMAx_Y_MODIFY (outer loop address increment registers), 9-18
DMEM_CONTROL (data memory control register), 6-24, 6-48
double-fault condition, 4-41
DPMC (dynamic power management controller), 8-2, 8-11
D port, 7-2
  interface, 7-8
DQM data I/O mask function, 16-27
DRxPRI SPORT input, 12-4
DRxSEC SPORT input, 12-4
DSP core interfaces, 7-5
DSP device ID register (DSPID), 18-26
DSPID (DSP device ID register), 18-26
DTEST_COMMAND (data test command register), 6-41
DTEST_DATAx (data test data registers), 6-42
DTxPRI SPORT output, 12-4
DTxSEC SPORT output, 12-4
dual 16-bit operations, 2-25
dynamic power management, 1-1, 8-1, 8-28, 8-29
dynamic power management controller (DPMC), 8-2, 8-11

E

EAB (external access bus), 7-2, 7-4, 7-20
  access arbitration, 7-21
  bandwidth, 7-22
  data transaction latency, 7-22
  EBIU, 16-3
early frame sync, 12-39
EBIU, 16-1
  access arbitration, 7-21
  arbitration to access, 7-21
  as slave, 16-3
  asynchronous interfaces supported, 16-1
  block diagram, 16-2

EBIU *(continued)*
bus error, 16-8
clock, 16-1
control registers, 16-7
external bus interface unit, 7-1
overview, 16-1
priority to access, 7-21
request priority, 16-1
status register, 16-7
EBIU_AMBCTL0 (asynchronous memory bank control 0 register), 16-12
EBIU_AMBCTL1 (asynchronous memory bank control 1 register), 16-12
EBIU_AMGCTL (asynchronous memory global control register), 16-9
EBIU_SDBCTL (SDRAM memory bank control register), 16-46
EBIU_SDGCTL (SDRAM memory global control register), 16-34
EBIU_SDRRC (SDRAM refresh rate control register), 16-50
EBIU_SDSTAT (SDRAM control status register), 16-49
EBUFE bit, 16-37
setting, 16-41
EMU core event, 4-19
emulation mode, 1-9, 3-9
enable interrupts (STI) instruction, 6-75, 8-19, 8-21
enabling interrupts, global, 4-39
endianess, 2-12
endian format, data and instruction storage, 6-66
entire field mode, PPI, 11-19
EPROM, 1-8
errors
bus parity, 4-49
bus timeout, 4-49
external memory addressing, 4-50
hardware, 4-49
errors *(continued)*
hardware conditions causing, 4-50
internal core, 4-49
misalignment of data, 6-72
multiple hardware, 4-49
peripheral, 4-49
system MMR, 4-50
error signals, SPI, 10-27
evaluation of loop conditions, 4-4
event
definition, 4-18
exception, 4-43
latency in servicing, 4-62
nested, 4-38
event controller, 3-1, 4-18
MMRs, 4-36
sequencer, 4-3
event handling, 1-9, 4-18
event monitor, 18-22
event processing, 4-3
event vector table (EVT), 4-40
EVT (event vector table), 4-40
EVX core event, 4-19
EX1 (execute 1), 4-7
EX2 (execute 2), 4-7
EX3 (execute 3), 4-7
EX4 (execute 4), 4-7
examples
SDRAM start address calculation, 16-54
exception, 4-1
events, 3-4, 4-43
events causing, 4-44
handler, executing, 4-48
multiple, 4-47
while exception handler executing, 4-49
exception handling, instructions in pipeline, 4-57
exception processing, deferring, 4-57
exception routine, example code, 4-60
exceptions, 1-10

exceptions by descending priority (table), 4-47
exclusive (definition), 6-75
EXCPT instruction, 4-48
execute 1 (EX1), 4-7
execute 2 (EX2), 4-7
execute 3 (EX3), 4-7
execute 4 (EX4), 4-7
execution cycle count registers (CYCLES and CYCLES2), 18-24
execution unit, components, 4-8
external access bus (EAB), 7-2, 7-4, 7-20
external bus (DEB), DMA, 7-4, 7-16, 7-17
external bus interface unit (EBIU), 1-13, 7-1
external bus interface unit. *See* EBIU
external clock oscillator, 17-5
external event watchdog (EXT_CLK) mode, 15-1, 15-39, 15-42, 15-47
external memory, 1-8, 6-47, 7-1
  design issues, 17-9
  interfaces, 16-5
external memory bus contention, how to avoid, 17-9
external memory controller, 1-12
external memory interfaces, 16-5
external memory map (figure), 16-4

F

Fast Fourier Transform, 2-35, 5-9
fetch address, 4-7
  incrementation, 4-8
fetched address, 4-2
FFT calculations, 5-9
FIFO, 16-1
FIO_BOTH (flag set on both edges register), 14-37
FIO_DIR (flag direction register), 14-8
FIO_EDGE (flag interrupt sensitivity register), 14-35
FIO_FLAG_C (flag clear register), 14-13
FIO_FLAG_D (flag data register), 14-12
FIO_FLAG_S (flag set register), 14-13
FIO_FLAG_T (flag toggle register), 14-17
FIO_INEN (flag input enable register), 14-39
FIO_MASKA_C (flag mask interrupt A clear register), 14-19, 14-22, 14-28
FIO_MASKA_D (flag mask interrupt A data register), 14-22, 14-28
FIO_MASKA_S (flag mask interrupt A set register), 14-19, 14-22, 14-28
FIO_MASKA_T (flag mask interrupt A toggle register), 14-22, 14-28
FIO_MASKB_C (flag mask interrupt B clear register), 14-19
FIO_MASKB_S (flag mask interrupt B set register), 14-19
FIO_POLAR (flag polarity register), 14-33
flag
  interrupt generation, 14-1
  transmit holding register empty status, 13-7
flag configuration register, core access to, 14-8
flag data register (FIO_FLAG_D), 14-12
flag direction register (FIO_DIR), 14-8
flag input enable register (FIO_INEN), 14-39
flag interrupt sensitivity register (FIO_EDGE), 14-35
flag mask interrupt A clear register (FIO_MASKA_C), 14-19
flag mask interrupt A data register (FIO_MASKA_D), 14-22, 14-28
flag mask interrupt A set register (FIO_MASKA_S), 14-19
flag mask interrupt A toggle register (FIO_MASKA_T), 14-22, 14-28

flag mask interrupt B clear register (FIO_MASKB_C), 14-19
flag mask interrupt B set register (FIO_MASKB_S), 14-19
flag mask interrupt registers, 14-19
flag polarity register (FIO_POLAR), 14-33
flags, programmable, 14-1
flag set on both edges register (FIO_BOTH), 14-37
flag toggle register (FIO_FLAG_T), 14-17
flag value registers, 14-10
flash, 7-1
flash memory, 1-8, 16-1
flex descriptor structure, 9-28
FLOW[2:0] field, 9-13
flushing buffers after semaphore release, 17-7
FLUSH instruction, 6-38
FLUSHINV instruction, 6-38
force interrupt / reset (RAISE) instruction, 3-11
fractional mode, 2-4, 2-14, C-5
fractional multiplier results format, 2-16
fractions, multiplication, 2-46
framed serial transfers, characteristics, 12-35
framed/unframed data, 12-34
frame pointer (FP), 4-4
frame pointer (FP) registers, 5-5
frame sync, 12-1
  active high/low, 12-37
  early, 12-39
  early/late, 12-39
  external/internal, 12-36
  internal, 12-30
  internally generated, 12-28
  late, 12-39
  multichannel mode, 12-51
  sampling edge, 12-37
frame sync *(continued)*
  setting up, 12-31
  SPORT options, 12-34
frame synchronization and SPORT, 12-3
frame sync pulse
  use of, 12-15
  when issued, 12-15
frame sync signal, control of, 12-14, 12-19
frequencies, clock and frame sync, 12-29
FSDR bit, 12-55
full duplex, 12-1, 12-4
full-on mode, 1-22

G

general-purpose interrupt, 4-18, 4-51
  with multiple peripheral interrupts, 4-34
general-purpose I/O (GPIO), 1-2
  pins, 14-1
general-purpose modes, PPI, 11-21
general-purpose timers, 1-1
global enabling and disabling interrupts, 4-39
ground plane, 17-13
ground planes, 17-15
GSM speech compression routines, 2-21
GSM speech vocoder algorithms, 2-42

H

H.100, 12-2, 12-55, 12-61
hardware conditions causing hardware error interrupts (table), 4-50
hardware-error interrupt (HWE), 4-49
  causes, 4-49
hardware errors, multiple, 4-49
hardware reset, 3-12
Harvard architecture, 6-4
heavy clock load and SDRAM, 16-39
hibernate state, 1-23, 8-15

high frequency design considerations, 17-11
high-speed measurements, 17-14
high-speed signals, capacitive loading on, 17-12
HMVIP, 12-61
hold, for EBIU asynchronous memory controller, 16-12
HWAIT (host wait) signal, SPI booting, 3-22
HWE (hardware-error interrupt), 4-49

I

I²S serial devices, 12-2
ICPLB address registers (ICPLB_ADDRx), 6-61
ICPLB_ADDRx (ICPLB address registers), 6-61
ICPLB data registers (ICPLB_DATAx), 6-56
ICPLB_DATAx (ICPLB data registers), 6-56
ICPLB fault address register (ICPLB_FAULT_ADDR), 6-64
ICPLB_FAULT_ADDR (ICPLB fault address register), 6-64
ICPLB_STATUS (ICPLB status register), 6-64
ICPLB status register (ICPLB_STATUS), 6-64
idle state, 3-9, 4-1
IF1 (instruction fetch 1), 4-7
IF2 (instruction fetch 2), 4-7
IF3 (instruction fetch 3), 4-7
I-fetch
  access exception, 4-47
  CPLB miss, 4-47
  misaligned access, 4-47
  multiple CPLB hits, 4-47
  protection violation, 4-47
ILAT (core interrupt latch register), 4-37
illegal combination, 4-47
illegal use protected resource, 4-47
IMASK (core interrupt mask register), 4-36
IMDMA, 1-2
IMDMA access request, 7-11
IMEM_CONTROL (instruction memory control register), 6-7, 6-48
index (definition), 6-76
indexed addressing, 5-9
  with immediate offset, 5-11
index registers (I[3:0]), 2-7, 5-2, 5-6
inner loop address increment registers (DMAx_X_MODIFY), 9-16 (MDMA_yy_X_MODIFY), 9-16
inner loop count registers (DMAx_X_COUNT), 9-16 (MDMA_yy_X_COUNT), 9-16
input clock (CLKIN), 8-1
input delay bit, 8-8
inputs and outputs, 2-25
inserting wait states using ARDY (figure), 16-21
instruction address, 4-3
instruction alignment unit, 4-7
instruction cache, coherency, 6-17
instruction decode (DEC), 4-7
instruction fetch 1 (IF1), 4-7
instruction fetch 2 (IF2), 4-7
instruction fetch 3 (IF3), 4-7
instruction fetches, 6-48
instruction fetch time loop, 4-17
instruction in pipeline when interrupt occurs, 4-57
instruction loop buffer, 4-17
instruction memory control register (IMEM_CONTROL), 6-7, 6-48
instruction memory unit, 4-7
instruction pipeline, 4-2, 4-6

instructions
ALU, 2-28, 2-30
conditional, 2-22
DAG, 5-16, 5-17
instruction set, 1-25
interlocked pipeline, 6-67
load/store, 6-67
multiplier, 2-39
protected, 3-4
register file, 2-8
shifter, 2-52
stored in memory, 6-66
synchronizing, 6-69
width, 4-8
instruction set, 1-25
instruction test command register (ITEST_COMMAND), 6-21
instruction test data registers (ITEST_DATAx), 6-22
instruction test registers, 6-20 to 6-24
instruction watchpoint address control register (WPIACTL), 18-7
instruction watchpoint address count registers (WPIACNTn), 18-5, 18-6
instruction watchpoint address registers (WPIAn), 18-5
instruction watchpoints, 18-4
instruction width, 4-8
integer mode, 2-14, C-5
integer multiplier results format, 2-17
integers, multiplication, 2-46
interface
D port, 7-8
on-chip L2 memory, 7-10
PAB bus, peripherals, 7-15
interfaces
DSP core, 7-5
external memory, 16-5
internal, 7-1
internal memory, 16-3
interface signals, asynchronous memory (table), 16-5
interface to 16-bit SRAM, 17-9
interface to 8-bit SRAM or flash, 17-2
interleaving
data in SPORT FIFO, 12-21
SPORT data, 12-5
internal address mapping (table), 16-47
internal bank, 16-28
internal clocks, 7-5
internal/external frame syncs. *See* frame sync
internal/external memory map, 1-7
internal interfaces, 7-1
internal memory, 1-8, 6-4
interfaces, 16-3
interrupt handling, instructions in pipeline, 4-57
interrupt priority register (IPRIO), 6-37
interrupts, 1-10, 4-1
buffer completion, 9-28
control of system, 4-18
definition, 4-18
DMA error, 9-28
enabling and disabling, 6-74
general-purpose, 4-18, 4-51
generated by peripheral, 4-20
global enabling and disabling, 4-39
hardware-error, 4-49
multiple sources, 4-22
nested, 4-38, 4-53
non-nested, 4-52
peripheral, 4-18, 9-28
PF pins, 14-1
priority watermark, 6-37
processing, 4-3
row completion, 9-28
servicing, 4-51
shared, 4-34
sources, peripheral, 4-27

interrupts *(continued)*
SPORT error, 12-27
SPORT RX, 12-25, 12-27
SPORT TX, 12-22, 12-27
timer, 15-3
interrupt service routine, determining source of interrupt, 4-28
interrupt status registers
(DMAx_IRQ_STATUS), 9-26
(MDMA_yy_IRQ_STATUS), 9-26
invalid cache line (definition), 6-76
I/O interface to peripheral serial device, 12-1
I/O memory space, 1-9
I/O pins, general-purpose, 14-1
IPEND (core interrupts pending register), 3-2, 4-38
IPRIO (interrupt priority register), 6-37
IRCLK bit, 12-19
IrDA, 13-14
receiver, 13-19
receiver pulse, 13-20
SIR protocol, 13-1
transmitter, 13-18
I-registers (index), 2-7
IRFS bit, 12-19, 12-36
ISR and multiple interrupt sources, 4-22
ITCLK bit, 12-14
ITEST_COMMAND (instruction test command register), 6-21
ITEST_DATAx (instruction test data registers), 6-22
ITEST registers, 6-21
ITFS bit, 12-15, 12-36, 12-52
ITU-656, 11-7, 11-14
IVG core events, 4-19
IVHW
core event, 4-19
interrupt, 4-49
IVTMR core event, 4-19

J

JTAG interface, 1-9
jump, 4-1
JUMP instructions, 4-9
conditional, 4-10
range, 4-11

L

L1 data memory, 1-8
L1 data SRAM, 6-26
L1 instruction memory, 1-8
L1 memory, on-core, 7-1
L1 memory. *See* level 1 (L1) memory; level 1 (L1) data memory; level 1 (L1) instruction memory
L1 or L2 memory, DMA access stalls, 7-13
L1 scratchpad RAM, 1-8
L1 SRAM
data memory, 7-1
instruction memory, 7-1
L2 bus, 7-2
L2 memory, 1-8
access bus arbitration, 7-12
access latency and throughput, 7-12
arbitration, 7-10
interface, 7-10
interface control logic clock rate, 7-10
L2 port
access priority, 7-11
access request, arbitration priority, 7-12
LARFS bit, 12-20
latched interrupt request, 4-37
late frame sync, 12-39, 12-50
latencies, 17-8
latency, 16-68
DMA data transfer, 7-13
EAB data transaction, 7-22
L2 memory access, 7-12
PAB transfer, 7-15

latency *(continued)*
programmable flags, 14-41
SDRAM read command, 16-59
servicing events, 4-62
setting CAS value, 16-42
when servicing interrupts, 4-51
LATFS bit, 12-16, 12-39
LB (loop bottom registers), 4-4
LC (loop counter registers), 4-4
least recently used algorithm (LRU), definition, 6-76
length registers (L[3:0]), 2-7, 5-2, 5-6
level 1 (L1) data memory, 3-11, 6-24
configuration, 6-5
sub-banks, 6-28
level 1 (L1) instruction memory, 3-11, 6-7
architecture, 6-11
configuration, 6-11
DAG reference exception, 6-9
dual-port capability, 6-10
instruction cache, 6-11
sub-bank organization, 6-7
level 1 (L1) memory
See also level 1 (L1) data memory; level 1 (L1) instruction memory
address alignment, 6-9
architecture, 6-4
definition, 6-76
L1 data SRAM, 6-5
L1 instruction SRAM, 6-5
scratchpad data SRAM, 6-6
level 2 (L2) memory, 6-44
latency, 6-45, 6-46
latency with cache off, 6-46
latency with cache on, 6-45
off-chip, 6-47
lines per frame register (PPI_FRAME), 11-13
line terminations, SPORT, 12-62
little endian (definition), 6-76
load, speculative execution, 6-70
load mode register command, 16-64
load operation, 6-67
load ordering, 6-68
load/store instructions, 5-5
logging nested interrupt, 4-56
logical operations, 2-23
logical shift (LSHIFT) instruction, 2-14, 2-48
logical shifts, 2-1
long jump (JUMP.L) instruction, 4-11
loop, 4-1
buffer, 4-17
conditions, evaluation, 4-4
disabling, 4-17
instruction fetch time, 4-17
registers, 4-4, 4-5
termination, 4-3
top and bottom addresses, 4-16
loop bottom registers (LB), 4-4
loop counter registers (LC), 4-4
loop top registers (LT), 4-4
low-power operation, 1-21
L-registers (length), 2-7
LRFS bit, 12-20, 12-35, 12-37
LTFS bit, 12-15, 12-35, 12-37
LT (loop top registers), 4-4

M

MACs (multiplier-accumulators), 2-36, 2-48
See also multiply without accumulate
dual operations, 2-47
multicycle 32-bit instruction, 2-46
managing core and system clocks, 17-5
mappings and priorities of interrupt events, 1-11
master host, booting, 3-21

master in slave out (MISO) pin, 10-3, 10-4, 10-5, 10-20, 10-22, 10-23, 10-26, 10-28
master out slave in (MOSI) pin, 10-3, 10-4, 10-5, 10-20, 10-22, 10-23, 10-26, 10-28
MCMEN bit, 12-50
MDMA_yy_CONFIG (DMA configuration registers), 9-12
MDMA_yy_CURR_ADDR (current address registers), 9-20
MDMA_yy_CURR_DESC_PTR (current descriptor pointer registers), 9-19
MDMA_yy_CURR_X_COUNT (current inner loop count registers), 9-21
MDMA_yy_CURR_Y_COUNT (current outer loop count registers), 9-22
MDMA_yy_IRQ_STATUS (interrupt status registers), 9-26
MDMA_yy_NEXT_DESC_PTR (next descriptor pointer registers), 9-10
MDMA_yy_PERIPHERAL_MAP (peripheral map registers), 9-23
MDMA_yy_START_ADDR (start address registers), 9-11
MDMA_yy_X_COUNT (inner loop count registers), 9-16
MDMA_yy_X_MODIFY (inner loop address increment registers), 9-16
MDMA_yy_Y_COUNT (outer loop count registers), 9-17
MDMA_yy_Y_MODIFY (outer loop address increment registers), 9-18
MemDMA
  channels, 9-41
memory
  See also cache; level 1 (L1) memory; level 1 (L1) data memory; level 1 (L1) instruction memory; level 2 (L2) memory
memory *(continued)*
  access, latency and throughput, 7-12
  address alignment, 5-13
  architecture, 1-6, 6-6
  asynchronous region, 16-2
  external, 1-8, 6-47, 7-1
  external memory, 16-5
  how instructions are stored, 6-66
  interface, 7-10
  internal, 1-8
  internal bank, 16-28
  L1 data, 6-24
  L1 data SRAM, 6-26
  level 2 (L2), 6-44
  management, 6-47
  moving data between SPORT, 12-42
  off-chip, 1-8
  off-core, 7-1
  on-core, 7-1
  page descriptor table, 6-50
  protected locations, 3-5
  protection and properties, 6-47
  terminology, 6-75
  transaction model, 6-66
memory architecture, 6-1
memory DMA, 9-40
  resgister naming conventions, 9-8
  transfer performance, 7-19
memory management unit (MMU), 6-47
memory map, 1-6
  external (figure), 16-4
memory-mapped registers (MMRs), 6-73, 6-75
  programmable flags, 14-8
memory page, 6-49
  attributes, 6-49
MFD field, 12-52
minimal clock load and SDRAM, 16-39
µ-law companding, 12-2, 12-33, 12-57
MMR location of core events, 4-40

MMR port, 7-2
mode
  boot, 1-24
  broadcast, 10-2, 10-14, 10-22
  SPI master, 10-2, 10-24
  SPI port master, 10-1
  SPI port slave, 10-1
  SPI slave, 10-2, 10-26
  UART DMA, 13-16
  UART non-DMA, 13-15
mode-fault error (MODF), 10-27
mode register, 16-28
modes
  multichannel, 12-46
  serial port, 12-10
MODF, 10-27
modified addressing, 5-4
modified (definition), 6-75
modify address, 5-1
modify instruction, 5-12
modify registers (M[3:0]), 2-7, 5-2, 5-6
moving data, serial port, 12-42
M-registers (modify), 2-7
MSEL clock frequency, 3-11
MSEL (multiplication select) field, 8-7
multichannel frame, 12-52
multichannel frame delay field, 12-52
multichannel mode, 12-46
  enable/disable, 12-50
  frame syncs, 12-51
  SPORT, 12-51
multichannel operation, SPORT, 12-46 to 12-61
multicycle bus lock, 17-7
multiple exception, for an instruction, 4-47
multiple interrupt sources, 4-22
multiple slave SPI systems, 10-14
multiplexed SDRAM addressing scheme (figure), 16-55
multiplier, 2-1
  accumulator result registers A[1: 0], 2-37, 2-38
  arithmetic integer modes formats, 2-15
  data types, 2-13
  fractional modes format, 2-15
  instruction options, 2-41
  instructions, 2-39
  operations, 2-37
  results, 2-38, 2-39, 2-42
  rounding, 2-38
  saturation, 2-39
  status, 2-22
  status bits, 2-39
multiplier select (MSEL) field, 8-3
multiply without accumulate, 2-44
MVIP-90, 12-61

N

NDSIZE[3:0] field, 9-14
nested interrupt, 4-38, 4-53
  logging, 4-56
nested interrupt handling (figure), 4-54
nested interrupts, 17-6
nested ISR
  example epilog code, 4-55
  example prolog code, 4-55
next descriptor pointer registers
  (DMAx_NEXT_DESC_PTR), 9-10
  (MDMA_yy_NEXT_DESC_PTR), 9-10
NMI, 4-43
  core event, 4-19
non-maskable interrupt (NMI), 1-10, 4-43
non-nested interrupt, 4-52, 17-6
non-nested interrupt handling (figure), 4-53
nonsequential program operation, 4-9
nonsequential program structures, 4-1
no operation (NOP) command, 16-67

NOP (no operation) command, 16-67
normal frame sync mode, 12-39
normal timing, serial port, 12-39
NOT, logical, 2-23
numbers
  binary, 2-3
  data formats, 2-11
  two's-complement, 2-4
  unsigned, 2-4
numeric formats, C-1, C-8
  binary multiplication, C-4
  block floating point, C-6
  integer mode, C-5
  two's-complement, C-1

O

off-chip memory, 1-8
on-chip memory, 1-8
open drain outputs, 10-23
open page, 16-26
operands for input, 2-37
operate CLKIN below specified frequency, 17-5
operating modes, 3-1, 8-12
  active, 1-22, 8-14
  deep sleep, 1-23, 8-14
  full-on, 1-22, 8-13
  hibernate state, 1-23, 8-15
  sleep, 8-14
OR, logical, 2-23
ordering
  loads and stores, 6-68
  weak and strong, 6-68
outer loop address increment registers
  (DMAx_Y_MODIFY), 9-18
  (MDMA_yy_Y_MODIFY), 9-18
outer loop count registers
  (DMAx_Y_COUNT), 9-17
  (MDMA_yy_Y_COUNT), 9-17
output, PF pin configured as, 14-1
output, PPI, 1 sync mode, 11-26
output delay bit, 8-8
overflow, data, 12-35
overflow, saturation of multiplier results, 2-39
overflow flags, 2-13

P

PAB (peripheral access bus), 7-2, 7-4, 7-14
  arbiter, 7-14
  arbitration, 7-14
  EBIU, 16-3
  errors generated by SPORT, 12-27
  peripherals interfacing to, 7-15
  transfer latency, 7-15
packing, serial port, 12-60
page size, 16-29, 16-48, 16-49
parallel peripheral interface (PPI). *See* PPI
parallel refresh command, 16-33
PC100 SDRAM standard, 16-1
PC133-compliant SDRAM controller, 1-8
PC133 SDRAM controller, 1-12
PC133 SDRAM standard, 16-1
PDWN bit, 8-9
performance
  DAB, DCB, and DEB, 7-18
  DAB bus, 7-18
  DCB bus, 7-18
  DEB bus, 7-20
  memory DMA transfers, 7-19
  PAB, 7-15
  programmable flags, 14-41
  SDRAM, 16-68
performance monitor control register (PFCTL), 18-20
performance monitor counter registers (PFCNTRn), 18-19
performance monitoring unit, 18-19

peripheral
configuring for an IVG priority, 4-35
DMA, 7-18
group 1, 7-1, 7-2, 7-4
group 2, 7-1, 7-4
interrupt generated by, 4-20
interrupt sources, 4-27
peripheral access bus (PAB), 7-2, 7-4, 7-14
peripheral clock (SCLK), 17-5
peripheral DMA start address registers, 9-11
peripheral interrupts, 4-18, 17-6
relative priority, 4-30
source masking, 4-29
peripheral map registers
(DMAx_PERIPHERAL_MAP), 9-23
(MDMA_yy_PERIPHERAL_MAP), 9-23
peripherals, 1-1
peripherals interfacing to PAB bus, 7-15
PFCNTRn (performance monitor counter registers), 18-19
PFCTL (performance monitor control register), 18-20
PF pins, shared with PPI, 11-1
PF. *See* programmable flags
PFX (programmable flag) pin, 3-22
pin
descriptions, 17-1
SPORT, 12-4
state during SDC commands (table), 16-62
terminations, SPORT, 12-62
pipeline
figure, 4-7
instruction, 4-2, 4-6
instructions when interrupt occurs, 4-57
interlocked, 6-67
pipelining, SDC supported, 16-41

PLL
active mode, 8-14
active mode, effect of programming for, 8-20
applying power to the PLL, 8-17
block diagram, 8-3
BYPASS bit, 8-14, 8-20
bypass mode, 3-11
CCLK derivation, 8-1, 8-3
changing CLKIN-to-CCLK multiplier, 8-17
clock counter, 8-11
clock dividers, 8-3
clock frequencies, changing, 8-11
clocking to SDRAM, 8-15
clock multiplier ratios, 8-3
code example, active mode to full-on mode, 8-21
code example, changing clock multiplier, 8-22
code example, full-on mode to active mode, 8-22
configuration, 8-3
control bits, 8-15
deep-sleep mode, effect of programming for, 8-20
disabled, 8-17
divide frequency (DF) bit, 8-3
DMA access, 8-13, 8-14, 8-21
dynamic power management controller (DPMC), 8-12
enabled, 8-17
enabled but bypassed, 8-14
full-on mode, effect of programming for, 8-20
lock counter, 8-11
maximum performance mode, 8-13
modification, activating changes to DF or MSEL, 8-19

PLL *(continued)*
modification in active mode, 8-16
multiplier select (MSEL) field, 8-3, 8-17
new multiplier ratio, 8-17
operating mode control bits, 8-18
operating modes, operational characteristics, 8-12
operating mode transitions, programming requirements, 8-18
operating mode transitions allowed, 8-18
operating mode transitions (figure), 8-18
PDWN bit, 8-15
PLL_LOCKED bit, 8-19
PLL_OFF bit, 8-17
PLL status (table), 8-12
power domains, 8-23
power savings by operating mode (table), 8-12
processing during PLL programming sequence, 8-20
relocking after changes, 8-19
removing power to the PLL, 8-17
RTC interrupt, 8-20
SCLK derivation, 8-1, 8-3
sleep mode, 8-14, 8-20
STOPCK bit, 8-15
voltage control, 8-12, 8-27
wake-up signal, 8-20
PLL control register (PLL_CTL), 8-6, 8-7
PLL_CTL (PLL control register), 8-6, 8-7
PLL divide register (PLL_DIV), 8-6
PLL_DIV (PLL divide register), 8-6
PLL_LOCKCNT (PLL lock count register), 8-6, 8-11
PLL lock count register (PLL_LOCKCNT), 8-11
PLL_OFF bit, 8-9
PLL programming sequence, 8-19
PLL_STAT (PLL status register), 8-6, 8-9
PLL status register (PLL_STAT), 8-6, 8-9
pointer register file, 2-5
pointer register modification, 5-12
pointer registers, 2-6, 3-4
point-to-point connections, design considerations, 17-12
polarity, programmable flags, 14-33
polling DMA registers, 9-54
popping, manual, 4-3
port access priority, 7-11
port access request, arbitration priority, 7-12
port connection, SPORT, 12-7
post-modify addressing, 5-1, 5-4, 5-7, 5-11
post-modify buffer access, 5-7
power dissipation, 8-23
power domains, 8-23
power-down warning, as NMI, 4-43
power management, 8-1, 8-28, 8-29
power up, 16-28
sequence, 16-64, 16-67
PPI, 11-1
active video only mode, 11-19
clock input, 11-1
control signal polarities, 11-3
data transfer, 11-33
data width, 11-3
DMA operation, 11-31
enabling, 11-9
entire field mode, 11-19
frame synchronization, 11-20
general-purpose modes, 11-21
MMRs, 11-3
operating mode, 11-3
output, 1 sync mode, 11-26
pins, 11-1
vertical blanking interval only mode, 11-19
video processing, 11-14
PPI_CONTROL (PPI control register), 11-3

PPI control register (PPI_CONTROL), 11-3
PPI_COUNT (transfer count register), 11-12
PPI_DELAY (delay count register), 11-11
PPI_FRAME (lines per frame register), 11-13
PPI_STATUS (PPI status register), 11-9
PPI status register (PPI_STATUS), 11-9
P port, 7-2
  interface, 7-7
precharge command, 16-29, 16-62
precharge delay, selecting, 16-44
PREFETCH instruction, 6-5, 6-38
P-register, 17-7
pre-modify addressing, 5-1
pre-modify instruction, 5-11
prioritization between the various internal buses, 17-8
priority
  EBIU access, 7-21
  L2 port access, 7-11
  Sys L2 port access, 7-12
processor mode
  determination, 3-1
  emulation, 3-9
  figure, 3-2
  identification, 3-2
  IPEND interrogation, 3-2
  supervisor, 3-6
  user, 3-3
processor state
  idle, 3-9
  reset, 3-10
processor state after reset, 17-1
product identification registers, 18-25
program counter register (PC), 4-2
  PC-relative indirect JUMP and CALL, 4-12
  PC-relative offset, 4-10, 4-11
program flow, 4-1
programmable flags, 1-2, 1-20, 14-1
  edge sensitive, 14-35
  latency, 14-41
  level sensitive, 14-35
  pins, interrupt, 14-1
  polarity, 14-33
  slave select, 10-12
  system MMRs, 14-8
  throughput, 14-41
programming information, xxxv
programming model
  cache memory, 6-4
  EBIU, 16-7
program sequencer, 1-4, 4-1
program structures, nonsequential, 4-1
protected instructions, 3-4
protected resources, 3-4
PSSE bit, 16-36
pulsewidth count and capture (WDTH_CAP) mode, 15-1, 15-30
pulsewidth modulation (PWMOUT) mode, 15-1, 15-20
pushing, manual, 4-3

Q

quad 16-bit operations, 1-4, 2-26
query semaphore, 17-7

R

range
  CALL instruction, 4-11
  conditional branches, 4-13
  JUMP instructions, 4-11
RCKFE bit, 12-20
RDIV
  equation for value, 16-51
  field, 16-50, 16-59
RDTYPE field, 12-19

read access, for EBIU asynchronous memory controller, 16-12
read transfers to SDRAM banks, 16-57
read/write command, 16-65
receive configuration registers (SPORTx_RCR1, SPORTx_RCR2), 12-16
receive FIFO, SPORT, 12-23
recommendations for unused pins, 17-1
reflection issues, 17-13
refresh, parallel, 16-33
register diagram conventions, xlv
register file instructions, 2-8
register files, 2-5, 2-10
register instructions, conditional branch, 4-10
register move, 4-13
registers
  accessible in user mode, 3-4
  core, A-1, A-13
  diagram conventions, xlv
  memory-mapped, A-1, A-13
  product identification, 18-25
  system, B-1
replacement policy, 6-34
  definition, 6-76
reserved SDRAM, 16-2
reset, 1-9
  core double-fault, 3-13
  core-only software, 3-13, 3-18, 3-20
  effect on memory configuration, 6-25
  hardware, 3-12, 8-15
  system software, 3-12, 3-16
  watchdog timer, 3-12, 3-16
reset initialization sequence
  programming for interrupts, 4-29
reset interrupt (RST), 4-41
RESET signal, 3-10
reset state, 3-10
resetting the processor, 17-1
reset vector, 4-42
reset vector addresses (table), 4-42
resources, protected, 3-4
resource sharing between multiple processors, 17-6
return address, 4-2
return address for CALL instruction, 4-10
return address registers, 4-3
return from emulation (RTE) instruction, 4-10
return from exception (RTX) instruction, 4-10
return from interrupt (RTI) instruction, 4-10
return from nonmaskable interrupt (RTN) instruction, 4-10
return from subroutine (RTS) instruction, 4-10
return instructions, 4-10
RETx register, 3-5
RFS pins, 12-34, 12-51
RFSR bit, 12-19, 12-34
ringing, to reduce, 17-12
RND_MOD bit, 2-18, 2-20
RO, xlv
ROM, 1-8, 16-1
rounding
  biased, 2-18, 2-20
  convergent, 2-18
  instructions, 2-17, 2-22
  unbiased, 2-18
round-to-nearest, 2-20
ROVF bit, 12-25, 12-26
row address, 16-48, 16-49
RRFST bit, 12-20
RSCLKx pins, 12-33
RSFSE bit, 12-20
RSPEN bit, 12-9, 12-16, 12-18
RST core event, 4-19
RST (reset interrupt), 4-41

RTE (return from emulation) instruction, 4-10
RTI instruction, use, 4-61
RTI (return from interrupt) instruction, 4-10
RTN (return from nonmaskable interrupt) instruction, 4-10
RTS (return from subroutine) instruction, 4-10
RTX (return from exception) instruction, 4-10
run length inductance, 17-12
RUVF bit, 12-25, 12-26
RX hold register, 12-24
RXSE bit, 12-20
RZI modulation, 13-18

S

SA10 pin, 16-34
sampling clock period, UART, 13-8
sampling edge, SPORT, 12-37
SCK1E bit, 16-38
SCLK[1] function and AMC, 16-39
SCLK (system clock), 8-4
  derivation, 8-1
  EBIU, 16-1
  frequency, 8-13
  status by operating mode (table), 8-12
scratchpad SRAM, 6-6
SCTLE bit, 16-34, 16-38
SDC, 16-3, 16-24, 16-49, 16-58
  commands, 16-61
  glueless interface features, 16-25
  operation, 16-58
  set up, 16-59
SDQM[1:0] encodings during writes for 16-bit SDRAM banks (table), 16-58
SDQM[3:0]
  (data I/O mask) pins, 16-33
  encodings during writes for 32-bit SDRAM banks (table), 16-57
  pins, 16-57
SDRAM, 7-1
  A10 pin, 16-34
  address mapping, 16-55
  auto-refresh, 16-65
  banks, 6-47, 16-29
  bank size, 16-1
  block diagram, 16-32
  buffering timing option (EBUFE), setting, 16-41
  clock enables, setting, 16-38
  configuration, 16-24, 17-9
  devices supported, 16-46
  example of start address calculation, 16-54
  external memory space, 6-2
  interface commands, 16-61
  interfaces, examples of, 17-9
  interface signals (table), 16-6
  memory banks, 16-4
  no operation command, 16-67
  operation parameters, initializing, 16-64
  performance, 16-68
  read command latency, 16-59
  read transfers, 16-57
  read/write, 16-65
  reserved, 16-2
  sizes supported, 6-47, 16-24
  start addresses, 16-1
  timing specifications, 16-67
SDRAM controller. *See* SDC
SDRAM control status register (EBIU_SDSTAT), 16-49
SDRAM memory bank control register (EBIU_SDBCTL), 16-46

SDRAM memory global control register (EBIU_SDGCTL), 16-34
SDRAM refresh rate control register (EBIU_SDRRC), 16-50
self-nested interrupts, 17-6
self-refresh command, 16-65
self-refresh mode, 16-29
  entering, 16-40
  exiting, 16-40
semaphore coherency, 17-7
semaphores, 17-6
  example code, 17-7
sensitivity, programmable flags, 14-35
SEQSTAT (sequencer status register), 4-4
sequencer registers, 3-4
sequencer status register (SEQSTAT), 4-3, 4-4
serial communications, 13-2
serial data transfer, 12-1
serial peripheral interface clock signal (SCK), 10-20, 10-22
serial peripheral interface slave select input signal (SPISS), 10-4, 10-13, 10-14, 10-20
serial peripheral interface (SPI) ports, 1-1, 1-16
serial peripheral interface (SPI). *See* SPI
serial ports. *See* SPORT
servicing interrupt, 4-51, 17-6
set associative (definition), 6-76
set (definition), 6-76
setting EBUFE, 16-41
set up
  for EBIU asynchronous memory controller, 16-12
  SDC, 16-59
  SDRAM clock enables, 16-38
shaded bits, xlvi
shared interrupt, 4-34, 4-56
shifter, 2-1, 2-48, 2-55
  data types, 2-14
  immediate shifts, 2-49, 2-50
  instructions, 2-52
  operations, 2-48
  register shifts, 2-50, 2-51
  status flags, 2-52
  three-operand shifts, 2-50
  two-operand shifts, 2-49
shifts, 2-1
short jump (JUMP.S) instruction, 4-11
SIC, 4-27
SIC_IAR0 (system interrupt assignment register 0), 4-34
SIC_IARx (system interrupt assignment registers), 4-30
SIC_IMASK (system interrupt mask register), 4-29
SIC_IWR (system interrupt wakeup enable register), 4-26
signal integrity, 17-12
signal integrity, suggestions to promote, 17-13
signed numbers, 2-3
sign extending data, 2-11
SIMD video ALU operations, 2-36
single 16-bit operations, 2-25
single step exception, 4-48
slave mode, booting, 3-21
slaves, EBIU, 16-3
slave select, SPI, 10-12
SLEN field, 12-14, 12-19
  restrictions, 12-32
  word length formula, 12-31
software interrupt handlers, 4-18
software management of DMA, 9-53
software reset modes, 17-1
software reset register (SWRST), 3-18
software watchdog timer, 1-19, 15-53
source termination, serial ports, 17-12

specifications, xxxv
speech compression routines, 2-21
SPI, 10-1
  beginning and ending transfers, 10-29
  block diagram, 10-2
  clock phase, 10-20, 10-22, 10-23
  clock polarity, 10-20, 10-23
  clock signal, 10-2, 10-23
  compatible peripherals, 10-1
  data transfer, 10-2
  detecting transfer complete, 10-15
  error signals, 10-27
  master mode, 10-2, 10-24
  multiple slave systems, 10-14
  point-to-point connections, 17-12
  registers, table, 10-19
  slave mode, 10-2, 10-26
  slave-select function, 10-12
  slave transfer preparation, 10-27
  switching between transmit and receive, 10-31
  transfer formats, 10-20
  transfer modes, 10-25
  transmission/reception errors, 10-15
  using DMA, 10-16, 10-32
SPI baud rate registers (SPI_BAUD), 10-7, 10-19
SPI_BAUD (SPI baud rate registers), 10-7, 10-19
SPI control registers (SPI_CTL), 10-4, 10-8, 10-19
SPI_CTL (SPI control registers), 10-4, 10-8, 10-19
SPI flag registers (SPI_FLG), 10-11, 10-19
SPI_FLG (SPI flag registers), 10-11, 10-19
SPI RDBR shadow registers (SPI_SHADOW), 10-18, 10-19
SPI_RDBR (SPI receive data buffer registers), 10-18, 10-19
SPI receive data buffer registers (SPI_RDBR), 10-18, 10-19
SPI (serial peripheral interface) ports, 1-1, 1-16
SPI_SHADOW (SPI RDBR shadow registers), 10-18, 10-19
SPI slave select, 10-12
$\overline{\text{SPISS}}$. *See* serial peripheral interface slave select input signal.
SPI_STAT (SPI status registers), 10-15, 10-19
SPI status registers (SPI_STAT), 10-15, 10-19
SPI_TDBR (SPI transmit data buffer registers), 10-17, 10-19
SPI transmit data buffer registers (SPI_TDBR), 10-17, 10-19
SPORT, 12-1 to 12-68
  active low vs. active high frame syncs, 12-37
  and DMA block transfers, 12-2
  channels, 12-46
  clock, 12-33
  clock frequency, 12-27, 12-29
  clock rate, 12-2
  clock rate restrictions, 12-31
  clock recovery control, 12-61
  companding, 12-33
  configuration, 12-10
  data formats, 12-32
  data word formats, 12-20
  delay when enabled, 12-10
  disabling, 12-10
  DMA data packing, 12-60
  enable/disable, 12-9
  enabling multichannel mode, 12-50
  framed serial transfers, 12-35
  framed vs. unframed, 12-34
  frame sync, 12-36, 12-39

SPORT *(continued)*
frame sync frequencies, 12-29
frame sync pulses, 12-1
framing signals, 12-34
general operation, 12-8
H.100 standard protocol, 12-61
initialization code, 12-19
internal memory access, 12-42
internal vs. external frame syncs, 12-36
late frame sync, 12-50
modes, 12-10
moving data to memory, 12-42
multichannel frame, 12-52
multichannel operation, 12-46 to 12-61
PAB error, 12-27
packing data, multichannel DMA, 12-60
pins, 12-1, 12-4
port connection, 12-7
receive and transmit functions, 12-1
receive clock signal, 12-33
receive FIFO, 12-23
receive word length, 12-24
register writes, 12-11
RX hold register, 12-24
sampling, 12-37
selecting bit order, 12-32
shortened active pulses, 12-10
single clock for both receive and transmit, 12-33
single word transfers, 12-42
source termination, 17-12
stereo serial connection, 12-8
stereo serial frame sync modes, 12-50
support for standard protocols, 12-61
termination, 12-62
timing, 12-62
transmit clock signal, 12-33
transmitter FIFO, 12-21
transmit word length, 12-21
TX hold register, 12-21

SPORT *(continued)*
TX interrupt, 12-22
unpacking data, multichannel DMA, 12-60
window offset, 12-54
word length, 12-31
SPORT error interrupt, 12-27
SPORT multichannel transmit select registers (SPORTx_MTCSn), 12-56
SPORT RX interrupt, 12-25, 12-27
SPORTs (serial ports), 1-1, 1-15
SPORT TX interrupt, 12-27
SPORTx_CHNL (SPORTx current channel register), 12-54
SPORTx current channel register (SPORTx_CHNL), 12-54
SPORTx_MCMCn (SPORTx multichannel configuration registers), 12-49
SPORTx_MRCSn (SPORTx multichannel receive select registers), 12-57
SPORTx_MTCSn (SPORTx multichannel transmit select registers), 12-58
SPORTx multichannel configuration registers (SPORTx_MCMCn), 12-49
SPORTx multichannel receive select registers (SPORTx_MRCSn), 12-56, 12-57
SPORTx multichannel transmit select registers (SPORTx_MTCSn), 12-56, 12-58
SPORTx_RCLKDIV (SPORTx receive serial clock divider registers), 12-27
SPORTx_RCR1 (receive configuration registers), 12-16
SPORTx_RCR2 (receive configuration registers), 12-16

SPORTx_RCR2 (SPORTx receive configuration register), 12-18
SPORTx receive configuration 2 registers (SPORTx_RCR2), 12-18
SPORTx receive data registers (SPORTx_RX), 12-23
SPORTx receive frame sync divider registers (SPORTx_RFSDIV), 12-28
SPORTx receive registers (SPORTx_RX), 12-51
SPORTx receive serial clock divider registers (SPORTx_RCLKDIV), 12-27
SPORTx_RFSDIV (SPORTx receive frame sync divider registers), 12-28
SPORTx_RX (SPORTx receive data registers), 12-23
SPORTx_STAT (SPORTx status registers), 12-25
SPORTx status registers (SPORTx_STAT), 12-25
SPORTx_TCLKDIV (SPORTx transmit serial clock divider registers), 12-27
SPORTx_TCR1 (transmit configuration registers), 12-11
SPORTx_TCR2 (transmit configuration registers), 12-11
SPORTx_TFSDIV (SPORTx transmit frame sync divider registers), 12-28
SPORTx transmit data registers (SPORTx_TX), 12-21
SPORTx transmit frame sync divider registers (SPORTx_TFSDIV), 12-28
SPORTx transmit registers (SPORTx_TX), 12-41, 12-51
SPORTx transmit serial clock divider registers (SPORTx_TCLKDIV), 12-27
SPORTx_TX (SPORTx transmit data registers), 12-21
SRAM, 1-8, 16-1
  interface, 17-9
  L1 data, 6-26
  L1 data memory, 6-5
  L1 instruction access, 6-9
  L1 instruction memory, 6-5
  scratchpad, 6-6
SRAM interfaces, examples of, 17-9
SRFS bit, 16-37
stack, 4-3
stack pointer (SP), 4-4
stack pointer (SP) registers, 5-5
stages of instruction pipeline (table), 4-7
stalls
  DMA access to L1 or L2 memory, 7-13
  pipeline, 6-67
start address calculation (table), 16-53
start address registers
  (DMAx_START_ADDR), 9-11
  (MDMA_yy_START_ADDR), 9-11
stereo serial data, 12-2
stereo serial device, SPORT connection, 12-8
stereo serial frame sync modes, 12-50
STI. *See* enable interrupts (STI)
STOPCK field, 8-9
store operation, 6-67
store ordering, 6-68
strong ordering requirement, 6-74
subroutines, 4-1
supervisor mode, 3-6
supervisor stack pointer, 5-5
support, technical or customer, xl
SWRST (software reset register), 3-18
SYNC, 6-70
synchronization of DMA, 9-53
synchronous DRAM (SDRAM), 1-8
synchronous serial data transfer, 12-1
SYSCFG (system configuration register), 4-6

SYSCR (system reset configuration register), 3-14
system
  DMA access request, 7-10
  L2 bus, 7-4
  overview, 7-7
  scratch pad SRAM memory, 7-1
system and core event mapping (table) events by priority (table), 4-19
system clock (SCLK), 8-1
system configuration register (SYSCFG), 4-6
system design, 17-1
  information, 17-1
  point-to-point connections, 17-12
  recommendations and suggestions, 17-13
  recommended reading, 17-15
system interrupt assignment register 0 (SIC_IAR0), 4-34
system interrupt assignment registers (SIC_IARx), 4-30
system interrupt controller (SIC), 1-10, 4-18
system interrupt mask register (SIC_IMASK), 4-29
system interrupts, 4-18
system interrupt wakeup enable register (SIC_IWR), 4-26
system-level interrupts, 17-6
system reset configuration register (SYSCR), 3-14
system software reset, 3-12, 3-16
system stack, recommendation for allocating, 4-62

T

tag (definition), 6-76
TBUFCTL (trace buffer control register), 18-16
TBUFSTAT (trace buffer status register), 18-17
TBUF (trace buffer register), 18-17
TCKFE bit, 12-16
TCKRE bit, 12-37
TCNTL (core timer control register), 15-50
TCOUNT (core timer count register), 15-51
TDM
  interfaces, 12-3
  multichannel mode, 12-2
TDTYPE bits, 12-14
technical support, xl
terminations, SPORT pin/line, 12-62
TESTSET instruction, 17-7
TFS pins, 12-41
TFSR bit, 12-15, 12-34
TFS signal, 12-52
throughput, 17-8
  achieved by interlocked pipeline, 6-67
  achieved by SRAM, 6-4
  L2 memory access, 7-12
  programmable flags, 14-41
  SPORT, 12-4
time-division-multiplexed (TDM) mode, 12-46
  See also SPORT, multichannel operation
timer
  auto-baud mode, 15-38
  external event watchdog (EXT_CLK) mode, 15-39, 15-42, 15-47
  interrupts, 15-3
  modes, 15-1
  pulsewidth count and capture (WDTH_CAP) mode, 15-30
  pulsewidth modulation (PWMOUT) mode, 15-20
  registers, 15-1
  watchdog, 15-53

timer input/output (TMRx) pin, 15-1
timer input select (TIN_SEL) bit, 15-38
timers, 1-17
  general-purpose, 1-1
  UART, 13-1
  watchdog, 1-2, 1-19
timer x configuration (T_CFGRx) registers, 15-2
timer x high word count (T_CNTHx) registers, 15-2
timer x high word period (T_PRDHx) registers, 15-2
timer x high word pulse width (T_WHRx) registers, 15-2
timing
  auto-refresh, 16-50
  SDRAM specifications, 16-67
timing examples, for SPORTs, 12-62
TLSBIT bit, 12-14
TOVF bit, 12-23, 12-26
TPERIOD (core timer period register), 15-52
trace buffer control register (TBUFCTL), 18-16
trace buffer exception, 4-48
trace buffer register (TBUF), 18-17
trace buffer status register (TBUFSTAT), 18-17
trace unit, 18-14, 18-17
traffic control/optimization, DMA, 7-13
transaction latency, EAB data, 7-22
transfer count register (PPI_COUNT), 11-12
transfer initiation from SPI master, 10-25
transfer latency
  DMA data, 7-13
  PAB, 7-15
transfers supported (table), 5-14
transmission format, SPORT, 12-2
transmit clock, serial (TSCLKx) pins, 12-33
transmit configuration registers (SPORTx_TCR1 and SPORTx_TCR2), 12-11
transmit frame sync (TFS) pins, 12-34
transmit holding register empty status flag, 13-7
t_{RAS}, 16-30
TRAS bits, 16-36, 16-44
t_{RC}, 16-31
t_{RCD}, 16-30
TRCD bits, 16-36
t_{RFC}, 16-31
TRFST bit, 12-16
t_{RP}, 16-30
TRP bits, 16-36, 16-45
TSCALE (core timer scale register), 15-53
TSFSE bit, 12-16
TSPEN bit, 12-9, 12-11, 12-13
TUVF bit, 12-22, 12-25, 12-41
t_{WR}, 16-31
TWR bits, 16-36, 16-45
TXF bit, 12-23, 12-25
TX hold register, 12-21
TXSE bit, 12-16
t_{XSR}, 16-31

U

UART
  autobaud detection, 15-38
  DMA channel latency requirement, 13-16
  DMA channels, 13-16
  DMA mode, 13-16
  IrDA receiver, 13-19
  IrDA receiver pulse, 13-20
  IrDA support, 13-18
  IrDA transmit pulse, 13-19
  IrDA transmitter, 13-18

UART *(continued)*
non-DMA mode, 13-15
sampling clock period, 13-8
standard, 13-1
timers, 13-1
universal asynchronous receiver transmitter ports, 1-1, 1-18
UART divisor latch register
(UART_DLH), 13-12
(UART_DLL), 13-12
UART_DLH (UART divisor latch register), 13-12
UART_DLL (UART divisor latch register), 13-12
UART_GCTL (UART global control register), 13-14
UART global control register (UART_GCTL), 13-14
UART_IER (UART interrupt enable register), 13-9
UART_IIR (UART interrupt identification register), 13-10
UART interrupt enable register (UART_IER), 13-9
UART interrupt identification register (UART_IIR), 13-10
UART_LCR (UART line control register), 13-4
UART line control register (UART_LCR), 13-4
UART line status register (UART_LSR), 13-6
UART_LSR (UART line status register), 13-6
UART ports, 1-18
UART_RBR (UART receive buffer register), 13-8
UART receive buffer register (UART_RBR), 13-8
UART scratch register (UART_SCR), 13-14
UART_SCR (UART scratch register), 13-14
UART_THR (UART transmit holding register), 13-7
UART transmit holding register (UART_THR), 13-7
UART transmit shift register (UART_TSR), 13-7
UART_TSR (UART transmit shift register), 13-7
unbiased rounding, 2-18
unconditional branches
branch latency, 4-14
branch target address, 4-14
undefined instruction, 4-47
underflow, data, 12-35
unframed/framed, serial data, 12-34
universal asynchronous receiver transmitter (UART) ports, 1-1, 1-18
unrecoverable event, 4-47
unsigned numbers, 2-4
data formats, 2-11
use of RTI instruction, 4-61
user mode
accessible registers, 3-3
entering, 3-5
leaving, 3-6
protected instructions, 3-4
user stack pointer (USP), 3-7, 5-5

V

valid bit
clearing, 6-38
figure, 6-24
function, 6-14
in cache-line replacement, 6-16
in instruction cache invalidation, 6-19
valid (definition), 6-76

vertical blanking interval only mode, PPI, 11-19
vias, 17-15
victim (definition), 6-77
video ALU instructions, 5-13
video ALU operations, 2-36
video instructions, 1-3
voltage, 8-23
  changing, 8-27
  dynamic control, 8-23
voltage-controlled oscillator (VCO), 8-3
voltage regulator control register (VRCTL), 8-24
VRCTL (voltage regulator control register), 8-24

W

W1C, xlv
W1S, xlv
wait states, additional, 16-19
WAKEUP signal, 3-10, 8-20
watchdog control register (WDOG_CTL), 15-56
watchdog count register (WDOG_CNT), 15-54
watchdog status register (WDOG_STAT), 15-55
watchdog timer, 15-53
  functionality, 1-19
watchdog timer reset, 3-12, 3-16
watchpoint match, 4-47
watchpoint status register (WPSTAT), 18-12
watchpoint unit, 18-1, 18-12
  combination of instruction and data watchpoints, 18-3
  data watchpoints, 18-10
  instruction watchpoints, 18-4
  memory-mapped registers, 18-2
  WPIACTL watchpoint ranges, 18-4
way
  1-way associative (direct-mapped), 6-75
  definition, 6-77
  L1 instruction memory as 4-way set-associative, 6-5
  priority in cache-line replacement, 6-16
WB (write back), 4-7
WDOG_CNT (watchdog count register), 15-54
WDOG_CTL (watchdog control register), 15-56
WDOG_STAT (watchdog status register), 15-55
WDSIZE[1:0] field, 9-15
WNR bit, 9-15
WOFF field, 12-54
word (definition), 2-6
word length
  SPORT, 12-31
  SPORT receive data, 12-24
  SPORT transmission, 12-2
  SPORT transmit data, 12-21
WPDACNTn (data watchpoint address count value registers), 18-11
WPDACTL (data watchpoint address control register), 18-12
WPDAn (data watchpoint address registers), 18-11
WPIACTL (instruction watchpoint address control register), 18-7
WPIAn (instruction watchpoint address registers), 18-5
WPSTAT (watchpoint status register), 18-12
wraparound buffer, 5-8
write, 6-77
write access for EBIU asynchronous memory controller, 16-12
write back (definition), 6-77
write back (WB), 4-7

write buffer depth, 6-37
write to precharge delay, selecting, 16-45
WSIZE field, 12-53

X

X16DE bit, 16-38
XOR, logical, 2-23

Y

YCrCb format, 11-5

Z

zero extending data, 2-11
zero-overhead loop registers, 4-4
zero-wait-state devices, faster switching of, 17-12